

THE IMPACT OF ARTIFICIAL INTELLIGENCE ON THE FUTURE OF WORKFORCES IN THE EUROPEAN UNION AND THE UNITED STATES OF AMERICA

An economic study prepared in response to the US-EU Trade and Technology Council Inaugural
Joint Statement

*The views expressed here do not necessarily represent the positions of the United States or
European Union or its Member States.*

Contents

| | |
|--|-----------|
| Introduction..... | 3 |
| a) Background for the Report..... | 3 |
| b) Scope of the Report..... | 3 |
| c) Executive Summary..... | 3 |
| Part I: Overview of AI..... | 4 |
| a) What Is AI? | 4 |
| b) Recent Progress on AI | 4 |
| c) Overall Progress and Future Directions | 7 |
| d) Economic Opportunities and Challenges Coming from AI | 8 |
| Part II: The Current State of AI Adoption..... | 11 |
| a) Adoption of AI in the United States | 11 |
| b) Adoption of AI in the European Union..... | 13 |
| Part III: The Impact of AI on Work..... | 15 |
| a) What Jobs and Worker Tasks Are at Risk from AI? | 17 |
| b) What New Jobs and Tasks Will Emerge from AI?..... | 19 |
| c) What Will Be the Impact of AI on Workers?..... | 20 |
| d) What Will Be the Impact of AI on the Workplace?..... | 22 |
| Part IV: Case Studies..... | 24 |
| a) Case Study 1: AI in Human Resources and Hiring | 24 |
| AI in Practice for Hiring | 24 |
| AI for Applicants | 26 |
| Algorithmic Creep and Unintended Consequences | 26 |
| Hiring and Job Loss | 28 |
| The Future of AI and Hiring | 29 |
| Conclusions..... | 29 |
| b) Case Study 2: Warehousing | 30 |
| Supply chains and the logistics industry | 30 |
| The growing importance of warehousing | 31 |
| Algorithmic management in warehousing | 37 |
| Working conditions in warehousing | 39 |
| Conclusions..... | 39 |
| Part V: Conclusions..... | 41 |
| References | 46 |

Introduction

a) Background for the Report

Both the US and European Commission (EC) expressed strong interest during the US-EU Trade and Technology Council in late September 2021 in working on a joint study to assess the potential impact of artificial intelligence (AI) on our workforces. The Pittsburgh statement committed to a joint “economic study examining the impact of AI on the future of our workforces, with attention to outcomes in employment, wages, and the dispersion of labor market opportunities. Through this collaborative effort, we intend to inform approaches to AI consistent with an inclusive economic policy that ensures the benefits of technological gains are broadly shared by workers across the wage scale” ([White House 2021](#), [European Commission 2021](#)).

b) Scope of the Report

Given the expansiveness of the possible scope of the project, this report is not designed to be exhaustive; rather, it highlights some of the most important themes for the economics of AI in a balanced manner. Because of the unique collaboration between the EC and the Council of Economic Advisers (CEA) on this work, our goal is to synthesize the perspectives of the US and European Union and academic work from both countries with a focus on implications relevant to policymakers. Our goal for this joint report is to strengthen collaboration on analysis and policy to ensure that the benefits of AI are broadly shared. The report is intended to highlight the economics behind AI-driven technological change with a particular focus on the institutional and policy decisions that will shape its future impact on the workforce.

c) Executive Summary

AI is a fast-evolving technology with great potential to make workers more productive, to make firms more efficient, and to spur innovations in new products and services. At the same time, AI can also be used to automate existing jobs and exacerbate inequality, and it can lead to discrimination against workers. While previous technological advances in automation have tended to affect “routine” tasks, AI has the potential to automate “nonroutine” tasks, exposing large new swaths of the workforce to potential disruption. The challenge for policymakers is to foster progress and innovation in AI while shielding workers and consumers from potential types of harm that could arise. Parts I and II of this report introduce AI and document its widespread adoption in the European Union and the United States; part III focuses on AI’s impact on labor; part IV contains case studies on hiring and logistics; part V concludes.

Part I: Overview of AI

a) What Is AI?

The OECD defines an AI system thus: “an AI system is a machine-based system that is capable of influencing the environment by producing an output (predictions, recommendations, or decisions) for a given set of objectives. It uses machine and/or human-based data and inputs to (i) perceive real and/or virtual environments; (ii) abstract these perceptions into models through analysis in an automated manner (e.g., with machine learning), or manually; and (iii) use model inference to formulate options for outcomes. AI systems are designed to operate with varying levels of autonomy” ([OECD 2019](#)).

As an illustration of the practical use of AI, consider asking AI to give its own definition of AI. The following prompt was submitted to a modern, trained AI engine called Generative Pre-Trained Transformer 3 ([GPT-3](#)): “write a paragraph describing what AI is and what it can do.” The output that followed from GPT-3 was:

AI stands for artificial intelligence. AI is the result of applying cognitive science techniques to artificially create something that performs tasks that only humans can perform, like reasoning, natural communication, and problem solving. AI is important because it has the potential to help humans become better at completing cognitive tasks, and to automate tasks that are currently difficult or impossible for humans to do.

– *GPT-3, a modern AI on the topic of AI*

This illustration hopefully highlights the transformative power of AI technology. While many previous technological advances led to the augmentation or automation of routine tasks, AI has the power to [augment or automate](#) a much wider variety of tasks that are normally thought to only be possible for humans to complete. In part III of this report, GPT-3’s thoughts on AI’s impact on the workforce are included.


b) Recent Progress on AI

The power of AI comes from its use of [machine learning](#), a branch of computational statistics that focuses on designing algorithms that can automatically and iteratively build analytical models from new data without explicitly programming the solution. It is a tool of prediction in the statistical sense, taking information you have and using it to fill in information you do not have. As shown in Figure 1, machine learning has been a dominant focus of AI research since the 1980s. Over the last 10 years or so, the uses of machine learning as a prediction technology have grown substantially. Machine learning is now commonplace: [Pandora](#) learns how to make better music recommendations based on its users’ preferences; [Google](#) can automatically translate content into

different languages based on translated documents found online; and [Facebook](#) predicts the responses of individuals to personalized adds to aid in the delivery of ads through its News Feed. One of the most common applications of machine learning is computer vision, or the use of computers to derive information from images and videos and is a major focus of research, reflecting its importance across a range of applications, from determining the content of images online for tagging or moderation, to enabling self-driving cars, to the retrieval of specific images or videos from databases.

Figure 1. AI research publications by topic, 1980-2021

Count of publications



Source: Microsoft Academic via [OECD:AI](#)

In the last half decade, there has been an increasing research focus on a specific subset of machine learning algorithms called [neural networks](#). These algorithms use a combination of weights and activation functions to translate a set of data inputs into predictions for outputs, measures the “closeness” of these predictions to reality, and then adjusts the weights it uses to narrow the distance between predictions and reality. In this way, a neural network can learn as it is fed more data. Networks with more than two layers of transformation between input and output are called “deep”. These architectures can learn hierarchical abstractions, which helps them efficiently

characterize complex relationships.

[Dean \(2019\)](#) summarizes the evolution of machine learning. Key ideas and algorithms underlying machine learning have been around since the 1960s. In the late 1980s and early 1990s, there was a surge of excitement in the AI community as people realized that machine learning could solve some problems in interesting ways, with substantial advantages stemming from their ability to accept raw forms of input data, and to train algorithms to perform predictive tasks. At that time, however, computers were not powerful enough to process vast amounts of data. It was not until the past several years, after decades of computational performance improvements driven by Moore’s Law, that computers finally started to become powerful enough to allow for this approach. Further, both public and private entities now have access to large and sophisticated data sets that can be used for the development and training of AI models. The availability of data—both in the form of physical exclusivity and in the form of formal intellectual property rights—can shape both the level and direction of innovative activity. This is explored by [Beraja, Yang, and Yuchtman \(2022\)](#), who show that Chinese firms with access to data-rich government contracts develop substantially more commercial AI software.

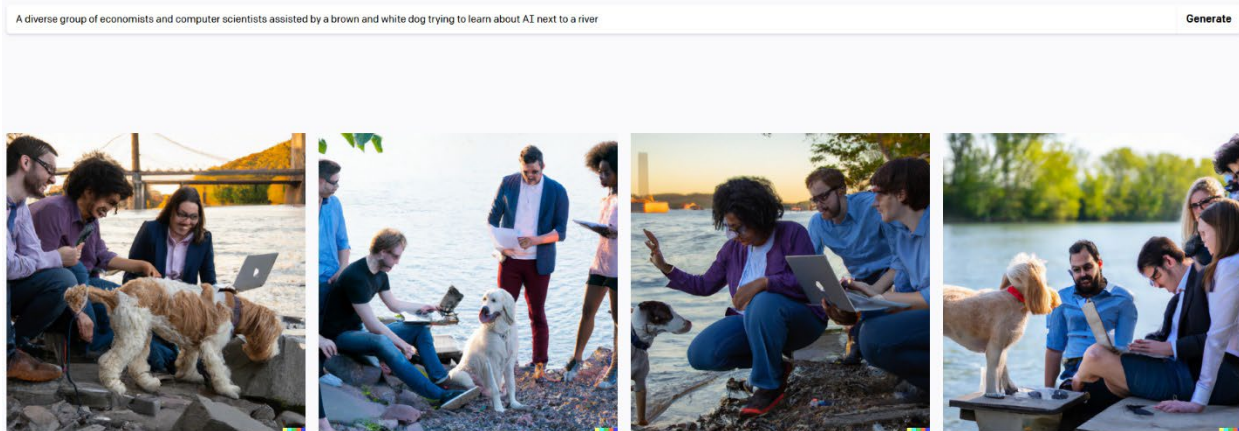
Consider a few examples of the progress that has been achieved using machine learning. First, Stanford University hosted the inaugural [ImageNet Challenge](#) in 2010. The challenge is, given a “training set” set of 1.2 million color images divided into 1,000 categories, to train a model to classify new color images to those same categories. The winning teams in 2010 and 2011 used traditional coding approaches and could not achieve error rates below 25 percent. In 2012, an entrant used a deep neural network for the first time and won with an error rate of 16.4 percent. Subsequent years saw innovations in deep learning applied to the problem with a winning error rate of only 2.3 percent in [2017](#), significantly lower than the average human tasked with the classification exercise ([Russakovsky et al. 2015](#)).

Second, consider [AlphaGo](#), a piece of software designed to play the ancient game Go against human players. It used neural networks and in addition to knowing the rules of Go, the model was trained both by playing against itself, as well as thousands of real amateur and professional games to learn strategies. In March 2016, AlphaGo beat the top-ranked player in the world 4 games to 1. Researchers then considered instead training the neural network by having it solely play games against itself—and the result was AlphaGo Zero. The neural network started with only random strategies and played against itself for 4.9 million games over three days. This new AI then defeated the previous version of AlphaGo by [100 games to 0](#).

Third, consider [DALL-E](#), which is based on the same technology as GPT-3. DALL-E is a model trained to generate images from a text description provided by a user. It was [trained](#) on a set of 250 million text–image pairs. The result is that it can create images that it has never “seen” but that fit the text description with which it was prompted. See Figure 2 for an example of the output given

when prompted with “a diverse group of economists and computer scientists assisted by a brown and white dog trying to learn about AI next to a river.”

Figure 2. An Example of AI Output



Source: [DALL-E](#)

These examples highlight the types of tasks that were previously thought to be impossible but now can be executed by AI, occasionally in a superior way to what a human could do.

c) Overall Progress and Future Directions

Progress in AI since the 1950s has been characterized by periodic cycles of breakthroughs and massive investment (“AI spring”) and periods of disappointment and little funding (“AI winter”). Technological breakthroughs lead to excited declarations of expected progress, which encourages increased investment. When the research stalls, “the enthusiasm, funding, and jobs would dry up” ([Mitchell 2021](#)). The 2010s were clearly a “spring,” with advances in image processing and natural-language processing as well as greatly increased computing power. Some have suggested that AI is now in a “Golden Age.” Nevertheless, there are concerns of a “winter” on the horizon given that some goals remain elusive, such as fully autonomous vehicles ([Mitchell 2021](#)). Computer scientists and philosophers thinking about the next leap forward in AI have emphasized the feasibility of a true artificial general intelligence (AGI) that equals or exceeds human intelligence. This AGI concept has been around since the era of electromechanical computing began after World War II. The first AI conference was held at Dartmouth College in 1956. In 1965, the Nobel laureate Herbert Simon predicted that “machines will be capable, within 20 years, of doing any work a man can do.” In recent years, AGI has seen a resurgence because of the development and advances in machine learning. While AGI is not a focal aspect of this study, the economic and societal impact of machines that surpass human intelligence would be extraordinary.

While the rise of AI promises both to improve existing goods and services and to greatly increase the efficiency with which they are produced, [Cockburn, Henderson, and Scott \(2019\)](#) argue that machine learning may have an even larger impact on the economy by serving as a new general-purpose technology (GPT) that is also an “invention in the method of invention” (IMI). What sets GPTs apart from IMIs is that IMIs that can also reshape the nature of the innovation process and the organization of research and development (R&D) itself. For example, Jumper et al. (2021) showed the successful use of their machine learning-based tool AlphaFold in predicting the physical structure of proteins and subsequently made available to the scientific community a [database](#) of over 200 million predicted protein shapes for researchers to use. Machine learning may be able to substantially “automate discovery” across many domains where classification and prediction tasks play an important role and may also expand the set of problems that can be feasibly addressed.

Previous IMIs help illustrate their importance. For example, the invention of optical lenses had an important direct economic impact on applications such as spectacles. But optical lenses in the form of microscopes, invented in the 17th century, also had enormous and long-lasting indirect effects on the progress of science: by making very small objects visible for the first time, microscopes opened the field of microbiology. Today, deep learning enables us to better understand genomes, thereby progressing the fields of molecular biology and genetics.

d) Economic Opportunities and Challenges Coming from AI

As AI technology continues to improve, it may have a substantial impact on the economy with respect to productivity, growth, inequality, market power, innovation, and employment. Policymakers could also use AI to create more efficient and equitable policymaking, as described in box 1 below.

Quantifying the benefits that AI will bring is difficult both because of the uncertainty of the future evolution of AI and also because the welfare contributions of AI—embedded in the proliferation of new and free goods such as search engines, digital assistants, or social media—are not captured well in our current national accounts. To this end, [Brynjolfsson et al. \(2019\)](#) propose a new metric called GDP-B, which quantifies their benefits rather than costs. Through a series of choice experiments, they estimate consumers’ willingness-to-pay for free digital goods and services. For example, including the welfare gains from Facebook would have added between 0.05 and 0.11 percentage points to GDP-B growth per year in the US. These are significant changes, especially considering that Facebook is just one product in the digital economy.

Box 1. Socially Optimal Applications of AI

Many policy problems facing governments require making decisions under uncertainty, but AI—combined with the large data sets to which governments have access—has given policymakers new tools to tackle that uncertainty. [Kleinberg et al. \(2018\)](#) use data from New York City on the decision by judges to either grant bail to a criminal defendant or require them to remain in prison. The authors find that—when compared with the decisions made by the judge—a machine learning algorithm could produce substantial welfare gains that could be allocated to either crime reduction (up to 24.7 percent) or to reduced incarceration (up to 41.9 percent), all while reducing racial disparities. However, the author’s note that while their model shows promise, actual implementation of such a prediction model as an aid to judges in their decision-making would require both a deep consideration of the objective courts are trying to achieve and how judges would actually use the model. In another case, [Aiken et al. \(2022\)](#) show that using machine learning algorithms and data from mobile phones improved the targeting of COVID-19 relief aid in Togo. In Togo, millions of dollars in aid were allocated to individuals during the early days of the COVID-19 pandemic based on a measure of need derived from machine learning models using mobile phone and satellite records. The authors found that, relative to an alternative proposal of allocating funds based on less granular measures of poverty, the algorithmic method increased the probability of getting the aid to the individuals who were most at risk.

However, there are real costs posed by AI to society that – as noted by [Acemoglu \(2021\)](#) – are all the more important to understand and confront because of “AI’s promising and wide-reaching potential”. Examples that stem directly from AI’s control of information include [privacy violations](#), creating [anti-competitive environments](#), and behavioral manipulation by machine learning techniques that enable companies to identify and exploit biases and vulnerabilities that consumers themselves do not recognize. Further, there is the direct risk that workers will be displaced by AI via excessive automation, as there is no guarantee that the current pace of the development of AI tools will achieve the socially optimal mix of automation and augmentation of tasks. Finally, there are a number of clear ways that AI have exacerbated social problems, including issues of discrimination and concerns about the functioning of democratic governments. There is substantial evidence—as discussed in Box 2 and in the hiring case study below— that AI has introduced and perpetuated [racial](#) or [other forms](#) of bias, both through issues with the underlying datasets used to make decisions, and by unintentional or seemingly-benign decisions made by algorithm designers. AI also can negatively impact how societies communicate on issues fundamental to the functioning of democracies, such as how echo chambers in social media can propagate false information and polarize society. While these costs are substantial, they are often not inherent to AI, but very much a product of the choices made in the development and deployment of the technology, meaning that there is a central role for governments in the studying, monitoring, and regulating of AI, as evidenced by the United States [AI Bill of Rights](#) and the

[European Commission Artificial Intelligence Act.](#)

The particular focus of this report is the impact of AI on the workplace. While in the four decades immediately after World War II, technological progress seemed to result in a prosperous labor market for all workers, a very different path of technological development started to emerge in the 1980s that was much less inclusive for low-paid workers, thus posing several challenges for policymakers. The resulting literature on skill biased technological change (SBTC), surveyed by [Acemoglu and Autor \(2011\)](#), documented how SBTC can account for trends in earnings distributions within the US and across different economics. However, the canonical SBTC model is one of technology generating a greater benefit to high versus low skill workers, while AI could either be a substitute or a complement to relatively higher skilled workers. One example is the strong incentives for firms to develop and adopt AI for automation instead of augmentation of tasks that were previously thought to require a human. Another example is that AI increases the ability to monitor workers. Although some amount of monitoring may be useful, monitoring can also be excessive if it is used to shift rents away from workers. In sum, unfettered AI could result in less democratic labor markets, worse working conditions, and an erosion of labor market institutions that favor workers.

Box 2. Failures in AI: Bias in Health Care

While AI algorithms have shown great promise in their ability to apply data to social and economic problems, there are many cases where they can exacerbate existing societal inequities. [Obermeyer et al. \(2019\)](#) examine the use of algorithms to determine which patients are “high-risk” and therefore receive additional resources and attention from care providers. The authors find that the algorithm they studied assigned poor patients to a lower risk score than an equally ill richer patient. This was because the algorithm used health cost as a proxy for health need, and poor patients generate lower costs than rich ones, potentially due to barriers in accessing care or bias they face in the health care system. The use of this proxy introduced bias into the algorithm and resulted in poor patients losing access to additional help they would otherwise have received.

AI’s enormous potential cannot be realized without a proper understanding and management of the above-noted challenges. [Mokyr \(2005\)](#) highlights the importance of this point by putting it in historical perspective. He argues that sustained economic growth after the onset of the Industrial Revolution was not solely due to the specific inventions made but was also facilitated by our understanding and management of these inventions. While there was also economic growth in the pre-1750 world from inventions such as gunpowder, spectacles, and the mechanical clock, this growth was not sustained because of a lack of understanding and management of these technologies. But how much does society understand about the consequences of the current era of AI and how can we best harness its enormous potential for sustained economic growth? Answering this question is the aim of this report, with a particular focus on the impact of AI on labor markets.

Part II: The Current State of AI Adoption

a) Adoption of AI in the United States


In the United States, the most recent publicly available data on the adoption of different technologies—including AI—come from the Census Bureau’s Annual Business Survey (ABS). Two recent papers—[Acemoglu et al. \(2022\)](#) and [McElheran et al. \(2022\)](#)—respectively use the 2019 and 2018 ABS modules to describe the adoption of AI technologies across US firms. Both papers find that overall adoption of AI is low, but that adoption is concentrated among a set of large, young firms. [McElheran et al. \(2022\)](#) focus on how owner and management characteristics correlate with AI adoption, finding that firms with younger, more educated, and more experienced owners are more likely to adopt AI technologies. [Acemoglu et al. \(2022\)](#) are able to take advantage of an expanded set of questions on AI adoption, and investigate the reasons behind firms’ adoption of AI, the barriers to further adoption, and the connection between AI adoption and productivity.

Both papers find that few firms overall have adopted AI, but that statistics about firm-level adoption mask the true share of US workers exposed to AI. [McElheran et al. \(2022\)](#) report that in 2017, 2.9 percent of firms used machine learning, 1.8 percent used machine vision, and 1.3 percent used natural-language processing. Similarly, [Acemoglu et al. \(2022\)](#) find that only 3.2 percent of US firms used AI as part of their processes and methods between 2016 and 2018. However, in 2017, 11.7 percent of workers worked at firms that used machine learning (and 6.8 and 8.8 percent were at firms that used machine vision and natural-language processing, respectively), and between 2016 and 2018, 12.6 percent of workers were employed at firms that utilized AI. This difference between firm and worker-level exposure stems from a key finding of both papers: larger firms are more likely to adopt AI technologies.

Important differences in AI adoption also exist irrespective of a firm’s size (see figure 3). First, firms in industries like information, professional services, management, and finance are the most likely to adopt AI technology. But workers in industries like retail trade, transportation, and utilities are also more likely to be exposed to AI than average. Second, irrespective of a firm’s size, younger firms are more likely to adopt AI. For example, of all large firms in the 95th to 99th percentiles of the firm size distribution, roughly 7 percent of firms in the youngest age quartile have adopted AI, whereas only about 3 to 4 percent of firms in oldest age quartile have done so. The fact that AI adoption is concentrated in larger and younger firms most likely reflects the fact that adopting this technology entails substantial costs and organizational barriers. Further, firms with venture-capital funding and other characteristics [McElheran et al. \(2022\)](#) categorize as “startup conditions consistent with high-growth entrepreneurship” are correlated with the use of AI.

Figure 3. Percentage of firms and workers with some AI Adoption

Percent (%)



Source: [Annual Business Survey 2019](#); CEA calculations

The 2019 ABS module also asks why firms adopt AI, and what barriers they face in implementing this technology. Both adopting and nonadopting firms report that the inapplicability of AI to the firm’s business, and AI being too costly, are the main reasons for either not adopting AI or the major existing barrier to further AI implementation. Of all AI adopters, around 80 percent (employment-weighted) report doing so to improve the quality of their product or service, 65 percent to upgrade existing processes, and 54 percent to automate existing processes. While the share of AI adopters stating that automation is one of their drivers is lower than other reasons, the findings of Acemoglu et al. (2022) regarding labor productivity—that AI adopters have higher labor productivity and lower labor shares than similar firms—are consistent with automation as a major application of AI. The use of AI to automate existing processes could have important adverse consequences for workers. AI competes more intensively with workers than other advanced technologies, with potentially important adverse consequences for the employment of individual workers.

The survey data discussed above are not a totally comprehensive look at firm-level adoption of AI technology. Many uses of AI might be missed in these surveys. For example, the use of voice assistance—like Siri, Cortana, or Alexa—is exceptionally common in the US. In 2017, 46 percent of Americans used digital voice assistance, according to the [Pew Research Center](#), with the vast majority using the service through their smartphone. In 2019, the share of Americans reporting the use of a digital voice assistant had grown to 72 percent, according to a [Microsoft](#) survey. This speaks to the emergence of AI in many areas of life, and not just the adoption of AI by firms.

The private sector is not the only part of the US economy that has begun to utilize AI. The US Federal Government has begun to implement AI in a range of [settings](#), including improving taxpayer [waiting times](#) when contacting the Internal Revenue Service (IRS) and creating AI competitions to predict [patient health outcomes](#) using Medicare data. The IRS, to address concerns about the long waiting times faced by callers, has implemented an AI-based voice bot system that currently allows taxpayers to set up payments and get notice questions answered. In the next year, this service will be expanded to allow for the bots to retrieve more information about individual taxpayers, further reducing waiting times. In 2019, the Centers for Medicare and Medicaid Services (CMS) created the CMS Artificial Intelligence Health Outcomes Challenge, a competition that was designed to accelerate “development of AI solutions for predicting patient health outcomes for Medicare beneficiaries.” In 2021, the competition [concluded](#), with the winners using Medicare case records to accurately predict patients who were likely to experience adverse events and explain these predictions to clinicians.

b) Adoption of AI in the European Union

The overall trends of firm-level AI adoption appear similar in the European Union to those in the United States: in 2021 8 percent of all enterprises with more than 10 employees employed AI technology. The [data](#), derived from Eurostat’s [Community Survey on ICT Usage and E-Commerce in Enterprises](#), asks about the usage of a range of AI technologies, including deep learning, the analysis of images and written/spoken language, and the automation of work. Larger firms were more likely to use some form of AI technology, with 28 percent of firms with more than 250 employees reporting their use. The survey also showed that firms most used AI to automate workflows, employ machine learning, or analyze written language (3 percent of firms in each case). The overall story is similar to that told by survey records from the [prior year](#): in 2020, 7 percent of enterprises in the EU reported to use AI. Some common uses of the technology were the analysis of large data sets via machine learning and the deployment of chatbots (2 percent of firms in both cases). With these data, we can also see the distribution of AI usage across EU Member States. In 2021, Denmark reported the largest share of enterprises employing AI, at 24 percent. Portugal (17 percent), Finland (16 percent), and Luxembourg and the Netherlands (both 13 percent) come next.

[Hoffman and Nurski \(2021\)](#) discuss the Eurostat Community Survey on ICT Usage and E-Commerce in Enterprises as well as other surveys (including a smaller [survey by the European Commission](#)) in more detail. As in the US, they find that robots are concentrated in manufacturing, while the adoption of other types of advanced technologies is higher in services such as finance, education, health, and social work. Within each of these sectors, larger firms are more likely to adopt AI, suggesting that there are substantial costs and organizational barriers involved in adopting AI. Skills and financial constraints are the leading reported barriers, with about 80 percent of enterprises citing a lack of skills in their internal workforce and in the external labor market, as well as the high cost of buying the technology and adapting their operational processes to AI.

Part III: The Impact of AI on Work

The previous parts of this report have argued that as AI continues to evolve and work its way into a wide variety of applications, its potential gains for society are enormous. As GPT-3 notes below in box 3, AI's benefits could span industries, providing workers with time for new tasks and firms with greater speed and accuracy through automation. The report has also shown how challenges arise for policy in seeking to benefit from the impact of AI—enumerated by GPT-3 in box 4 –, including policies related to the collection and control of data, communication in society, and safeguarding democratic principles.

This part of the report focuses on the opportunities and challenges resulting from AI for labor markets. AI has the potential to increase productivity, create new jobs, and raise living standards. However, by its very nature of performing “non-routine” tasks formerly thought to be strictly the domain of humans, AI is likely to disrupt large swaths of jobs and tasks. This may lead to difficult adjustments for workers as jobs are redesigned or required skills change.

Box 3. AI Prompt: “Write a paragraph describing how AI can benefit the workforce”

In Part I, GPT-3 (a modern, trained AI engine) was asked to give its own definition of AI. Next consider the following prompt submitted to [GPT-3](#): “write a paragraph describing how AI can benefit the workforce.” The output that followed from GPT-3 was:

“AI has the potential to streamline tasks across many industries, from health care to manufacturing. For example, in health care, AI can be used to automate the tasks of sorting through medical images to diagnose conditions. In manufacturing, AI can be used to automate the tasks of quality control and inspection. By automating these tasks, AI has the potential to free up workers' time so that they can focus on higher-level tasks. In addition, AI has the potential to improve the accuracy of these tasks, as well as to identify tasks that are most suitable for automation.”

– *GPT-3, a modern AI on the benefits of AI for workers*

This poses several challenges for policymakers. One example illustrated in the previous part of the report is that much of the development and adoption of AI is intended to automate work instead of augmenting it. Private firms advancing AI technology are likely to do so in a direction that maximizes profits, which may not be the socially desirable direction. Another example is that AI increases the ability of employers to monitor workers. In sum, while the potential benefits of AI for labor markets are numerous, unfettered AI could also result in a less democratic and less fair labor market.

Box 4. AI Prompt: “What are the downsides of AI for workers?”

Consider the following prompt submitted to [GPT-3](#): “What are the downsides are of AI for workers.” The output that followed from GPT-3 was:

- “1. Job Losses: One of the potential downsides of AI is that it could lead to increased unemployment as machines begin to replace human workers in a variety of industries.
2. Inequality: Another potential downside of AI is that it could exacerbate existing economic inequality as the benefits of AI technology disproportionately accrue to those who are already wealthy and have access to the best resources.
3. Security Risks: Another potential downside of AI is that it could create new security risks as malicious actors begin to use AI technology for malicious purposes.
4. Ethical Concerns: Another potential downside of AI is that it could raise ethical concerns as AI technology increasingly begins to impact our human lives in ways that we may not be comfortable with.”

– *GPT-3, a modern AI on the downsides of AI for workers*

Although the literature identifies more challenges than those listed by GPT-3 above, designing policies to tackle these challenges requires a proper understanding of how technological progress has affected labor markets in the recent past, and how AI is likely to change jobs in the future. [Autor \(2022\)](#) provides an overview of the recent thinking about the impact of digital technologies on labor markets. His starting point is a task-based view of labor markets that has become the standard framework in the literature over the past decade.

The hypothesis arising from this view is that digital technologies can automate “routine tasks.” What makes a task routine is that it follows an explicit, fully specified set of rules and procedures. Tasks fitting this description can be codified in computer software and executed by machines (e.g., robots to assemble a car, email to deliver messages). Conversely, “non-routine tasks” have historically been challenging to program because the explicit steps for accomplishing these tasks are often not formally described. Paradoxically, even though we cannot formally express non-routine tasks in an algorithm, many of these tasks are easy for humans to do. This is known as Polanyi’s paradox—that “humans know more than they can tell,” named after the 20th-century philosopher Michael Polanyi and his argument that all our knowing is rooted in tacit knowledge.

[Goos, Manning, and Salomons \(2014\)](#) show that routine tasks are concentrated in middle-paid occupations (e.g., machine operators, office clerks), while non-routine tasks (e.g., waiting tables in a restaurant, cleaning a room, diagnosis diseases, or team management) are concentrated in low-paid occupations (e.g. restaurant server, cleaner) and high-paid occupations (e.g., health professionals, managers). Consequently, typical automation technologies have decreased demand

for middle relative to low-paid and high-paid occupations, resulting in a process of job polarization. They show that this is happening in the 16 Western European countries that they examine from 1993 to 2010, and similar evidence exists for the US ([Acemoglu and Autor 2011](#)).

AI has the potential to fundamentally change the relationship between automation technology, labor demand, and inequality. While studies have so far examined digital technologies such as computers and industrial robots, AI overturns the assumption that technology can accomplish only routine tasks. AI can be used to infer tacit relationships that cannot be fully specified by underlying software, because AI learns to perform these tasks inductively by training on examples instead of by following explicit rules that are programmable.

Consequently, many non-routine tasks done in both low-paid and high-paid occupations that cannot be performed by computers could be performed by AI in the future, with very different implications for labor demand, job polarization, and inequality. For example, we might no longer see a process of job polarization but one of stronger relative employment growth in high-paid occupations (if AI automates non-routine tasks in low-paid occupations) or of stronger relative employment growth in low-paid occupations (if AI automates non-routine tasks in high-paid occupations).

Because of AI's promise of a paradigm shift in our thinking about its impact on work and inequality, there is much uncertainty about AI's implications for labor markets. The remainder of this part of the report focuses on these four questions:

- a) What jobs and tasks are at risk from AI?
- b) What new jobs and tasks will emerge from AI?
- c) What will be the impact of AI on workers?
- d) What will be the impact of AI on the workplace?

a) [What Jobs and Worker Tasks Are at Risk from AI?](#)

Although earlier digital technologies automated occupations that were intensive in doing routine tasks (e.g., machine operators, office clerks), AI as a prediction technology has the potential of also automating various non-routine tasks across a wide range of occupations. To study this question, a small but rapidly growing literature has emerged that applies a task approach to analyze the effects of AI adoption on different occupations ([Acemoglu et al. 2022](#); [Brynjolfsson, Mitchell, and Rock 2018](#); [Felten, Raj, and Seamans 2020](#); [Webb 2020](#)). These studies do not start from the premise that AI can only do a given set of tasks. Instead, they rely on various innovative ways to determine what worker tasks AI can and cannot automate.

[Webb \(2020\)](#) offers one example. He uses natural-language processing (NLP) algorithms that exploit the overlap between the text of job task descriptions and the text of patents to develop a new method for identifying which tasks can be automated by any technology. This allows him to construct a measure of the “exposure” of occupations to that technology. For example, suppose a doctor’s job description includes the task “diagnose patient’s condition.” An NLP algorithm then extracts the verb–noun pairs from this task, which would be “diagnose condition.” The algorithm then quantifies the same verb–noun pairs in a different sample of patents to identify whether any technology could automate a doctor’s tasks.

Using this approach, Webb (2020) first examines the impact of two previous types of new technologies: software and robots. For software, exposure is decreasing with education, with individuals in middle-wage occupations most exposed. Men are much more exposed to software than women, reflecting the fact that women have historically clustered in occupations requiring complex interpersonal interaction tasks, which software is not capable of performing. For robots, individuals with less than a high school education, and men under the age of 30 years are most exposed. By and large, these results are consistent with the literature on job polarization, which has found that computers and robots reduced demand for routine, middle-wage jobs while increasing it for non-routine, low- and high-wage jobs between 1980 and 2010.

Webb’s (2020) study then turns to the impact of AI on the demand for occupations. In contrast to software and robots, AI performs tasks that involve detecting patterns, making judgments, and optimizing. The most-exposed occupations include clinical laboratory technicians, chemical engineers, optometrists, and power plant operators. More generally, high-skill occupations are most exposed to AI. Moreover, as might be expected from the fact that AI-exposed jobs are predominantly those involving high levels of education and accumulated experience, it is older workers who are most exposed to AI. There are also some low-skilled jobs that are highly exposed to AI. For example, production jobs that involve inspection and quality control are exposed. However, these constitute only a small proportion of low-skill jobs.

To conclude, an emerging body of research suggests that AI can outperform workers in an increasing set of complex tasks mainly done by educated workers. Compared with earlier digital innovations, this suggests a paradigm shift in our thinking about AI’s potential to automate worker tasks. For example, the automation of worker tasks by AI could exacerbate a process of occupational deskilling instead of job polarization.

This paradigm shift will not be straightforward. One reason commonly argued is that, because AI is unaware of the rich context of many real-world problems, it cannot accomplish the complex-decision tasks that humans regularly undertake in their work. Another reason is pointed out by [Acemoglu et al. \(2022\)](#): that so far, AI has no detectable effects on the labor market at the aggregate

occupation level. This current absence of any visible aggregate effects of AI could lower our sense of urgency to understand its impact on work, even when such effects appear likely in the future.

b) What New Jobs and Tasks Will Emerge from AI?

In capturing AI's benefits, an important lever for policymakers is that AI not only automates but also augments work. History is full of examples of jobs that were predicted to be doomed by automation but that instead flourished and were transformed. The [introduction of the first ATMs](#) around 1970 was predicted to end the job of traditional bank tellers, but the US today instead has many more bank tellers, at many more bank branches, doing different tasks than before because ATMs are poorly suited to such as relationship banking ([Bessen 2015](#)). If the set of tasks were fixed, then advancing automation would crowd workers into an ever-narrowing subset of tasks, perhaps finally making human labor altogether obsolete, if AI would evolve into a state of AGI. However, it is possible that even AGI will create many new jobs for workers.

While AI's potential to automate jobs has received relatively little attention, even less is known about AI's potential to create new jobs for workers. However, it is possible to learn from the larger literature that asks how many new jobs does technological progress create? To answer this question, [Autor et al. \(2022\)](#) exploit the emergence of new job titles in the US Census Bureau's occupational descriptions that survey respondents supply on their Census forms. Their analyses show that, irrespective of whether a new job is created because of technological progress or some other reason, new work is quantitatively important. They estimate that more than 60 percent of US employment in 2018 was found in job titles that did not exist in 1940. Examples of new titles are "fingernail technician," which was added in 2000, and "solar photovoltaic electrician," which was added in 2018. Interestingly, "artificial intelligence specialist" first appeared in 2000.

Turning to the nature of new work, they find that between 1940 and 1980, most new work that employed non-college workers was found in middle-skilled occupations. After 1980, however, the locus of new work creation for non-college workers shifted away from these middle-tier occupations and toward traditionally lower-paid personal services. Conversely, new work creation employing college-educated workers became increasingly concentrated in professional, technical, and managerial occupations. In combination, these patterns indicate that new work creation has polarized after 1980, mirroring (and in part driving) aggregate job polarization.

To further explain the creation of new job titles, and the role of technological progress, Autor et al. (2022) follow a procedure like Webb (2020) by examining patent data using NLP. Different from Webb (2020), however, is that they also instructed their NLP algorithm to look for text that indicates augmentation instead of automation of worker tasks. For example, in 1999, the U.S. Patent and Trademark Office granted a patent for a "method of strengthening and repairing fingernails." Their algorithm links this patent to the occupational title of "Technician, fingernail," which was added by Census Bureau in 2000. Similarly, their algorithm links the 2014 patent

“systems for highly efficient solar power conversion” to the occupational title of “Solar photovoltaic electrician,” which was added in 2018. In sum, Autor et al. (2022) show that new technologies are an important driver of the creation of new worker tasks.

Autor et al. (2022) also instruct their NLP algorithm to look for text in patents that indicate a new technology’s potential to automate (instead of augment) worker tasks. What they find is that some occupations, such as radiologic technologists and machinists, have a high rate of automation relative to augmentation. Therefore, labor demand and thus employment would tend to fall in these occupations. Conversely, in other occupations, including industrial engineers and analysts, augmentation has been more important than automation, resulting in an increase of employment in these occupations. Interestingly, many occupations are either simultaneously exposed to both augmentation and automation or are not exposed to any technology at all. Examples of occupations with very limited exposure to technological progress as of yet include jobs that require interpersonal skills such as childcare workers, hotel clerks, and clergy.

To conclude, though technology’s potential to automate jobs has received widespread attention, it also augments work and is an important driver of new job creation. Autor et al. (2022) term this double-sided impact of innovation on work “the race between automation and augmentation.” In occupations with declining (increasing) employment shares, this race is won by automation (augmentation). Understanding this race gives policymakers important levers to capture the benefits of technological progress. For example, a race between automation and augmentation of worker tasks, even within narrowly defined occupations, implies that new technologies can perhaps be steered toward more augmentation and less automation.

Autor et al. (2022) do not specifically focus on AI. But many new jobs augmented by AI may soon enter as new occupational titles—digital assistant engineer, warehouse robot engineer; and content-tagger on social media, among other jobs. An important policy question is whether these are the jobs that society wants AI to create. Simultaneously, employment in many high-paid jobs that is desirable from a policy perspective might be eroded by AI’s potential to automate their tasks.

c) What Will Be the Impact of AI on Workers?

The impact of technological progress, including AI, on work is characterized by competing forces of automation and augmentation of worker tasks, even (and mainly) within narrowly defined occupations. The focus of researchers—as well as managers, entrepreneurs, and policymakers—should therefore be not only on AI’s automation or augmentation potential but also on job redesign. For example, [Brynjolfsson, Mitchell, and Rock \(2018\)](#) conjecture that machine learning will require a substantial redesign of tasks for concierges, credit authorizers, and brokerage clerks. The need for job redesign also poses challenges for worker adaptability: worker skills to do certain tasks, and worker mobility across jobs in the labor market.

Worker Skills

[Acemoglu et al. \(2022\)](#) leverage a new module introduced in the US Census Bureau's 2019 ABS not only to assess firms' adoption of AI but also to explore firms' self-assessment on the implications of AI for their demand for labor and skills. Among AI adopters, 15 percent report that AI increased overall employment levels and 6 percent indicate that AI decreased them, which points to the limited and somewhat ambiguous effects of AI on employment levels. Instead, 41 percent of AI adopters increased their skill demand, while almost no firms (less than 2 percent) report a reduction in their demand for skills. This self-reported increase in firms' skill requirements when they adopt AI explains part of the well-known skills gap and highlights the importance of investments in worker skills.

[Genz et al. \(2022\)](#) provide similar evidence for Germany. They examine how German workers adjust to firms' investments into new digital technologies, including AI, augmented reality, or 3D printing. For this, they collected novel data that link survey information on firms' technology adoption to administrative social security data for Germany. They then compare technology adopters with non-adopters. Though they find little evidence that AI affected the number of jobs, the absence of an overall employment effect masks substantial heterogeneity across workers. They find that workers with vocational training benefit more than workers with a college degree. One explanation might be that AI augments vocational workers more than it augments tasks done by college workers. Another explanation is that Germany's traditionally strong vocational training system (76 percent of all workers in the sample completed vocational education) provides an abundance of specialized skills that direct the development and adoption of AI toward making use of (and thereby augmenting) vocational skills.

Worker Mobility Across Jobs

It is inevitable that workers in some jobs will be displaced because AI automates rather than augments worker tasks and/or workers no longer have the required skills to do their jobs. Job displacement is costly for those made redundant and could be disruptive for labor markets in general. These adjustment costs and disruptions were also characteristic of previous technological upheavals, exemplified by the automation of the role of telephone operator as discussed in box 5. However, because of the rapidity with which AI is evolving, these costs now may be particularly acute, but research documenting the transition of displaced workers to new jobs (or not) due to AI is very limited.

One exception is [Bessen et al. \(2022\)](#). Using Dutch administrative data, they examine what happens to workers who are made redundant when their firm invests in AI with the purpose of automating the firm's production process. They find that the expected annual income loss across all workers before their firm adopts AI accumulates to 9 percent of one year's earnings after 5

years. They also show that this annual income loss is driven by spells of unemployment within a year (rather than, e.g., quickly moving into lower-paid jobs), with unemployment benefits only insuring partially against their income losses. These adverse effects of AI automation are larger in smaller firms, and for older and middle-educated workers. In sum, their results suggest that there are substantial adjustment costs for displaced workers, and that these costs are only partially offset by unemployment insurance. Relatedly, in a case study given in part IV below, AI's role in the hiring process is highlighted. In some ways, AI can improve the transition between jobs by facilitating matches between employers and employees, although there are also potential drawbacks discussed in that setting.

Box 5. Labor Market Adjustment after the Automation of Telephone Operating

[Feigenbaum and Gross \(2022\)](#) examine the introduction of mechanical switching in operating telephone calls that took place in half of all US States between 1920 and 1940. They study adjustments in the labor market for young female telephone operators, one of women's main occupations at the time. They find that telephone operators were significantly less likely to still be working as operators 10 years after their state's cutover to mechanical switching. While some found other jobs in the telephone industry, others (especially older workers) left the workforce, and those who remained employed were more likely to have switched to lower-paying occupations. They also find that automation of telephone operating did not decrease overall demand for young women in their local labor markets. After automation, young women were less likely to become telephone operators, and they entered different jobs such as middle-skilled clerical and lower-skilled service occupations (mainly typists and waitresses).

d) What Will Be the Impact of AI on the Workplace?

AI will also drastically change how we design our workplaces and companies' business models. In turn, these changes will affect working conditions.

[Wood \(2021\)](#) discusses the prevalence of algorithmic management of workplaces. Algorithmic management relies on data collection and surveillance of workers to manage workforces in an automated way. Online labor platforms are a well-known example. These platforms enable workers to choose the clients and jobs they take, how they carry out those jobs, and the rates they charge to do them. However, to varying degrees, workers' ability to make these choices is strongly shaped by platform rules and design features. Increasingly, algorithmic management is also being used in other settings, such as in warehouses, retail, manufacturing, marketing, consultancy, banking, hotels, call centers, and among journalists, lawyers, and the police. Wood (2021) summarizes several detailed case studies from these sectors.

Consider the case of digital platforms for taxi services or home deliveries. On these platforms, algorithms allocate tasks to drivers via their smartphones (or other handheld devices). For example, a taxi platform can notify a driver with a trip request, which the driver must accept within a 15-second window. Only after having accepted the request, the algorithm provides drivers with the passenger's location, fare, and destination. The limited time frame given by the algorithm to accept a request while withholding key information is done to minimize the chances of drivers' declining trip requests. Moreover, if drivers decline too many requests, the algorithm can log them temporarily out of the app as a punishment. Once a driver has accepted a trip request, the algorithm recommends a route for reaching the drop-off location. If drivers deviate from the suggested route, the algorithm can send notifications. If the app is also responsible for paying drivers, the app can further punish drivers who take too long to get to their destinations by refusing to release drivers' payments. In sum, despite the many advantages platforms bring to workers and their clients, their algorithmic management can strongly reduce workers' ability to choose clients, how to do their tasks, and the rates they charge to do them.

Weil (2017) discusses the broader impact that algorithmic management has on business models and labor relations. In his [testimony](#) to the U.S. House of Representatives, he argues that firms can use information and communication technologies to erode the need for traditional employment relationships. Since the 1980s, many large corporations have shed their role as direct employers, in favor of outsourcing work to smaller subcontractors or franchisees. Competition between these subcontractors or franchisees implies that costs, including wages, are lower compared with a situation where the lead corporation directly employs these outsourced workers. Because this fissuring of workplaces, as Weil calls it, mainly affects low-wage jobs, it has exacerbated higher wage inequality, decreased occupational safety, and increased health risks for workers in fissured jobs. Unfettered AI can become the glue to make the overall business strategy of fissuring operate even more effectively. It can further enable lead companies and their shareholders to manage their labor supply chains even better through the intelligent monitoring of outsourced workers.

Part IV: Case Studies

a) Case Study 1: AI in Human Resources and Hiring

The traditional approach to hiring during the latter half of the 20th century was straightforward: applicants submitted their résumé and cover letter, perhaps with answers to job-specific questions, to opportunities they saw advertised either on job boards or in classified ads. Hiring managers waded through a stack of applicant files, winnowing through the list to determine who they wanted to consider. After a series of interviews, a job offer was made, and the particulars of the offer were perhaps negotiated. Eventually, a candidate would accept an offer and begin work.

During the past decade of AI innovations across modern economies, the hiring process has dramatically changed. While the individual steps of the process are broadly similar, at each stage firms have adopted AI-based tools to increase the speed and scale of the process. AI can match résumés with job listings on a massive scale, saving both the applicant and hiring manager much time. AI can screen résumés to discard applicants that are likely to be a poor fit; it can then put the candidate through assessments to further narrow the field. For many firms, only at the later stages of the process do humans enter the picture: final interviews, negotiations, and convincing a candidate to accept an offer remain important tasks for HR professionals. However, once the candidate has accepted, AI returns to assist in the retention and promotion roles. While the central goal of hiring remains the same, the set of tools available has changed, primarily due to innovations in AI that go far beyond the suite of tools previously available to hiring managers.

In order to explore the recent developments of AI in the field of hiring, staff at the Council of Economic Advisors conducted a series of interviews with stakeholders in this space. Over the span of the summer of 2022, they conducted six interviews with four firms, representatives of an industry group, and a scholar in the field of AI. Each was asked a series of questions about the current use of AI in hiring, with an emphasis on fostering an open and honest discussion. The interviews were compiled by CEA staff and combined with independent research and consultations with European Commission partners to form this case study.

AI in Practice for Hiring

Consider a firm that is looking to hire for a number of open jobs. They want to find the right candidates for each position as fast as possible, a task that requires both maximizing match quality as well as speed. Further, this requires high-volume hiring, as the firm is now facing the same rising job turnover to which many firms in the [modern labor market](#) are exposed. This is new for the firm, as historically it was not posting as many jobs and did not feel the pressure to hire with much speed. This new world of hiring is quantitatively different in scale, with an increase in job postings, applications submitted, and offers extended. This pressure to process more applications faster and attract more diverse and qualified workers without sacrificing the quality of a given

match, leads firms to turn to AI solutions. The hiring manager is now faced with managing his or her team of recruiters and juggling multiple job openings, each of which is at a different stage in the recruitment process. However, at each stage they can turn to AI for guidance, advice, and support. Trey Causey of the job search site Indeed offered some observations about the broader hiring landscape; “it is hard to think of a place in hiring where AI is not appearing. Hiring managers are able to couple AI solutions with human participation in the hiring process.”

At the very beginning of the process, the hiring manager needs to post a job opening, including crafting the text of the job description that will be published across a range of job search platforms. However, they do not need to do this alone. Rather, they can turn to a range of services that will use natural-language processing to help them write job descriptions. The power of these tools is that they link language to a data set of outcomes, allowing the hiring manager to [craft job descriptions](#) that will maximize the chances of attracting the right kind of applicants.

Turning from writing the text for a new job posting, the hiring manager now needs to figure out how to get this opportunity in front of candidates. To do this, they make use of one of the most common applications of algorithms in hiring: the matching of job applicants with job postings. These algorithms rely on the text of résumés and job postings, as well as information about the positions and the background of the candidates to determine which candidates are the best fit for a given job posting, or vice versa. In some cases, this results in a literal “fit score,” which hiring managers can use in evaluating candidates. The use of these systems may require the hiring manager to oversee the purchase of advertising on different recruitment platforms in order to get the job posting in front of the right candidates.

A few days ago, the hiring manager posted a separate job opening, and already candidates are reaching out for details about the position and with questions about the application process. However, the hiring manager is not responding to these messages. Rather, a small army of [chatbots](#)—powered by natural-language processing—are tasked with responding to specific and unique questions from candidates about the open positions. And the use of chatbots does not end there. The hiring manager can turn to chatbots to screen the initial round of applicants, an important step given the volume of applications the firm is receiving. These bots collect information on the candidates’ backgrounds and experience that will be incorporated into the decision of whether to advance a candidate to the next round.

With the new, smaller pool of applicants, the hiring manager can now utilize a range of evaluation tools, ranging from recorded interviews that are transcribed and analyzed to “gamified” assessments, essentially logic puzzles and other games that can assess particular skills among applicants. As a result, they can assess a candidate’s personality, skills, and critical thinking, all without interviewing a candidate directly. And these tests have science to back them up, connecting the results of the test to the narrow skill set they are measuring. However, the hiring

manager is cautious about using these tools, both because the connections between the skill sets and job performance have not been as thoroughly reviewed, and also because they have seen older tools be withdrawn because of issues of blatant bias. At the same time, the hiring manager finds these tools useful, as they both increase the speed of the application process and have the potential to improve the quality of matches.

AI for Applicants

In preparing for the job hunt, all applicants—whether applying directly from college, transitioning between roles within an industry, or thinking about a change of careers—can turn to a number of AI-based tools for developing their interview skills and preparing their résumé. Firms such as Indeed and VMock offer AI-powered tools that often highly weight specific terms to assess one’s résumé and offer suggestions on ways to improve it. In particular, because many résumés now are screened by AI-based tools, an important way to improve a résumé is to use keywords that help a candidate survive an initial screening.

Another way AI can help applicants is by focusing on applicant skills that may apply to a broader set of potential jobs than the candidate had considered. One example given by VMock was that the day-to-day work of a chef involves managing a large number of people in a high-pressure environment while meeting tight deadlines, a set of skills valuable in a wide number of positions outside food service. ZipRecruiter’s platform uses an active learning algorithm to try to understand which open positions appeal most to a candidate based on their interest level in the positions they have shown them so far; they use a similar algorithm on the hiring side to learn what types of candidates hiring managers are looking for. These learning algorithms lead to better matches on both the applicant and hiring sides.

After a candidate has applied to a position, they may be interacting with some of the chatbots described in the previous section; natural-language processing technology has advanced to the point that chatbots provide human-like interactions with job candidates. The firm Indeed [candidly discusses](#) the pros and cons of chatbots, noting that while they can be useful in reducing unconscious human biases in the hiring process, they also have the potential to create a negative impression with candidates. This reflects a larger question: to what degree are candidates aware of bots when they are being evaluated by a human versus an algorithm?

Algorithmic Creep and Unintended Consequences

At almost every stage the hiring process at firms across the United States and Europe, the role of AI-driven algorithms is increasing. This trend—dubbed “algorithmic creep” by Alex Engler of the Brookings Institution—encompasses both the expanded use of algorithms across different stages of the hiring process and the higher share of firms employing algorithms at each stage. Most of the firms that were spoken to for this report claimed that the result of this widespread adoption is a

process they claim is faster, scales up dramatically, and results in a greater number of qualified candidates finding better jobs. However, as Engler notes, “this transition to an algorithm-dominated hiring process is happening faster than firms, individuals, or governments are able to evaluate its effects.”

One of the primary concerns raised by nearly everyone interviewed is that greater adoption of AI-driven algorithms could potentially introduce bias across nearly every stage of the hiring process. Machine learning algorithms are often referred to as “[money laundering for bias](#),” in that they give the appearance of a fair and clean mathematical process while still exhibiting biases. Some firms are aware of this risk and are conscious of the potential for AI to lead to a less-biased hiring process than the former human-centric hiring process was known to be. However, even fair and well-intentioned algorithms have been shown to introduce bias in unexpected ways. For example, [Lambrecht and Tucker \(2019\)](#) show that STEM career ads that were explicitly meant to be gender neutral were disproportionately displayed by an algorithm to potential male applicants because the cost of advertising to younger female applicants is higher and the algorithm optimized cost-efficiency. A team of researchers at Google [studied](#) how natural-language models interpret discussions of disabilities and mental illness and found that various sentiment models penalized such discussions, creating bias against even positive phrases such as “I will fight for people with mental illness.”

The AI-powered tools developed to evaluate applicants have been particularly troublesome in this respect. The Center for Democracy & Technology issued a complete report in December 2020 titled “[Algorithm-Driven Hiring Tools: Innovative Recruitment or Expedited Disability Discrimination?](#)” The report explored the challenges algorithmic evaluation of candidates face in complying with the Americans with Disabilities Act, noting the many ways in which different screening tools may be biased against those with disabilities. Because many AI assessment tools employ a video interview, it is further worth noting that a study of automated speech recognition software [found large racial disparities](#) between interviews with white and African American individuals. Further, the [Gender Shades](#) project by the MIT Media Lab shows that three leading AI tools perform systematically worse at analyzing images of darker-skinned individuals, and particularly darker-skinned women, indicating that algorithmic tools introduce error for particular groups. These studies raise serious questions about the introduction of bias when employing AI in candidate evaluation and firms [are responsible](#) if they implement AI solutions that violate existing laws and regulations on discrimination.

There was friction in most of our conversations between the knowledge that the pre-AI hiring process was far from bias-free—there is evidence of long-standing bias against [nonwhite workers](#), and [people with disabilities](#), among other groups—and the concern that the use of AI can exacerbate these biases. However, in each conversation, there was—to a lesser or greater extent—the belief that correctly applied AI could reduce bias in hiring. Engler stated that the transition to

the heavy utilization of AI in hiring has the opportunity to “functionally be a re-set when it comes to labor-market discrimination against workers across racial, gender, disability, and economic lines.” And the Data and Trust Alliance stressed that they were working on utilizing AI to identify existing bias within models, and to help their clients meet their individual diversity and inclusion goals. [Audits of AI technology](#) are increasingly seen as necessary step in implementing new AI systems, although clarifying exactly how these audits should be conducted is still an evolving discussion within the field. There are also ongoing efforts within the private sector to create best practices. The [Data & Trust Alliance](#), for example, is a consortium of firms that has developed a set of safeguards in the form of a list of questions that firms can ask of their AI tool vendors to evaluate the potential for bias in their algorithms. The degree to which these efforts have begun to move us toward a more egalitarian world of hiring is very much unknown, leaving all the actors in this space split between the possibilities and perils of AI and bias.

Hiring and Job Loss

One concern that emerges with any discussion of automation is the potential for job loss, in this case within the hiring and human resources profession. As discussed earlier in this report, AI has the capacity to automate non-routine tasks, and the discussions with firms highlighted how, if an HR department is using an AI algorithm to schedule appointments, review résumés, answer candidates’ questions, and, overall, automate a range of tasks—from the mundane to the intricate—it is conceivable that they will need fewer (or different) workers.

The consequence of this change in roles is that HR departments may be looking for a different combination of skills or experiences from their HR professionals. An example raised by Causey of Indeed is that HR managers now need to understand how to manage the promotion of their job listing through various internal and external tools utilizing AI. Though many job-posting platforms allow free posting services, most offer the opportunity to “promote” a job to increase its exposure to more candidates. AI poses the opportunity to make this simple and reduce demands on hiring managers. Different sites have different models for this: some use a “cost-per-click” model, where the firm pays every time the job posting is clicked, while others use a “cost-per-application” model, where the firm pays every time it receives an application. Both types of systems require hiring managers to set daily budgets for how much to spend. This task is new to the HR profession and requires learning about the different systems as well as the value of investing in promoting job posts.

However, some firms we interviewed made the point that this is a potential benefit of the adoption of AI in hiring. Mahe Bayireddi, CEO and cofounder of Phenom, stressed that many firms with slim profit margins are looking to AI as a way to both increase the efficiency of their hiring and reduce the costs associated with HR. He framed Phenom’s work as helping firms ask “where can you personalize, and where can you automate.” With this framing, he both highlighted how AI can

allow firms to connect with more workers at a greater speed while preserving humans for activities where more strategic value is needed, like convincing someone to accept an offer once it has been extended. This reframing of the HR role as managers of AI and as “talent advisers” makes certain elements of the hiring process more human.

The discussions above mirrored the questions posed in Part III regarding the difficulty of ascertaining the net effect of AI on employment. The AI processes discussed by firms may end up automating many of the tasks currently performed by people, although AI may make employees more productive and create new tasks that require human intervention. While the net effect of AI on employment is unclear, it is likely to result in an HR office that is able to manage a much larger scale of operations.

The Future of AI and Hiring

The firms consulted for this report were asked how AI had changed the hiring process, and many noted that AI has *quantitatively* changed the process by allowing the large-scale implementation of systems to attract, screen, and assess potential employees. However, *qualitatively*, the HR function has not changed; it is still the search for the right person for each job opening. AI has been applied as a tool to the existing hiring process; there is optimism, however, that some structural changes in the hiring process may be coming.

One such change cited by several firms was the possibility of “[digital credentials](#),” or “[learning and employment records](#),” as technologies that could improve how AI functions in the roles it plays today. Standardized electronic records for education, training, and skills could streamline how AI matches applicants with job postings, potentially in a fairer and more balanced way. Of course, this relies on the assumption that the process of obtaining such records is itself fair and unbiased. Though AI has proven effective at using existing résumés and job listings, it is not clear that those are the best “inputs” to a matching system. As a result, firms have been thinking about what might be the right way to design the process from scratch. There is therefore great potential for more transformation of the HR process from AI—Indeed’s Causey noted that “we are on the cusp of a lot of qualitative changes” to hiring.

Conclusions

The overarching message from discussions with firms in the hiring space was that AI-powered algorithms could improve nearly every step in the hiring process for firms, HR staff members, and candidates. In fact, some firms very clearly structured their responses to questions by systematically addressing each stakeholder in the hiring process, discussing how each could benefit from the greater deployment of everything from chatbots to predictive models that match candidates with potential employers. However, adoption of AI has been so rapid that firms may

not entirely recognize the implications of allowing algorithms into the HR function. Firms should be auditing their usage of AI tools to ensure compliance both with labor regulations and also with their own ethical standards.

b) Case Study 2: Warehousing

Supply chains and the logistics industry

In the 1990s, supply chain management became a major field as more and more industries reassessed how products could be made because of new information technologies, combined with the reduction of many international trade restrictions such as quotas and tariffs. Supply chains benefit companies through gains from specialization, resulting in lower prices for intermediate goods and services. Consumers benefit through lower prices for final goods, and workers benefit through increased employment and wages. At the core of supply chain operations are logistics and warehousing to move goods between companies and to consumers.

Supply chains lengthen when companies decide to outsource activities because the benefits of finding outside suppliers capable of providing intermediate inputs at lower cost outweigh the benefits of keeping those activities inside the organization. Supply chains often take the shape of deeply integrated networks with lead companies at the center and their suppliers orbiting around them, resulting in business models known as lean manufacturing, lean retailing, and global-value chains.

Lean manufacturing

Lean manufacturing is a core production strategy famously developed by Toyota after WW-II. Its objective is to reduce inventories of intermediate parts and finished products, carefully matching real-time demand for goods with the quantity moving through the supply. In most supply chains, this requires high levels of coordination at each step in the process, careful management of capital and labor, attention to quality and factors that affect throughput, and strong logistics support systems. Lean manufacturing started in the auto industry, and many other manufacturing and retail sectors have adopted some or all of Toyota's pioneering practices.

Lean retailing

Like lean manufacturing, [lean retailing](#) takes advantage of information technologies, automation, industry standards, and innovations in logistics and warehousing to align orders from suppliers more closely with what consumers are buying in the store. By using sales information collected through millions of scans of bar-coded labels, retailers reduce their need to stockpile large inventories of products, thereby reducing their risks of stock-outs, markdowns, and inventory carrying costs.

Global-value chains

The reduction in quotas and tariffs together with [falling](#) costs of transportation transformed domestic supply chains into global-value chains. In manufacturing, this implies seeking suppliers for parts and assemblies that are outsourced to producers abroad (instead of at home). Outsourcing to producers abroad is also known as offshoring.

Logistics and warehousing – the core of supply chains

As lean manufacturing, lean retailing, and global-value chains spread across sectors in the economy in the past decades, the importance of logistics and warehousing increased. Moreover, the emergence of digital technologies since the 1980s and, more recently AI, have transformed the nature of logistics and warehousing.


Initially, a warehouse was simply the place where you store inventory—and where that inventory can sit for long periods of time. Although warehousing required tracking and managing where things had been left, it didn't require a lot of attention to how quickly those things could be accessed and moved once needed. However, with lean production, the warehouse becomes a distribution center—a place where intermediate or final products are efficiently tracked, processed, and moved. Some modern distribution centers are also known as “fulfillment centers” or “FC”.

The growing importance of warehousing

Driven by the emergence of supply chains, the economic importance of warehousing as a sector has increased over the past decades. This is illustrated in Figure 4, using data from the OECD's Structural Analysis (STAN) database, which includes the US and European countries with sufficiently detailed data. For each country, the first black bar shows the percentage of a country's total value added in 2018 that is produced by the NAICS Revision 4 subsector “Warehousing and support activities for transportation”. Countries in Figure 4 are ranked by their importance of warehousing in total value added in 2018. In Ireland, warehousing accounted for 0.5 percent of total value added. In the US, this number was 0.7 percent. The highest numbers are 2.7 percent for Belgium and 3.6 percent for Lithuania. For each country, the second blue bar shows the percentage of a country's total workforce in 2018 that is employed in the subsector of warehousing. For example, the highest fraction is found for Belgium, with 2.3 percent of the total workforce employed in warehousing. In all countries, the subsector of warehousing is an important employer.

Figure 4: Warehousing value-added and employment shares by country

Warehousing as a percent of countries total (%)



Source: [OECD STAN database](#) for industrial analysis, 2020 ed.

Notes: Warehouse industry defined as the NAICS Rev. 4 subsector “D52: Warehousing and related transport support activities”

The markers in Figure 4 also show that warehousing has become more important relative to other sectors since 1995. For example, in Germany – the EU’s largest economy – the share of warehousing in total value-added increased from 1.0 percent in 1995 to 1.8 percent in 2018 (with value added expressed in 2015 euros in both 1995 and 2018). In the US, the value-added share of warehousing was 0.4 percent in 1995 and 0.7 percent in 2018. Overall, Figure 4 shows a growing importance of warehousing in advanced economies since 1995, in line with the rapid expansion of supply chains.

To see what drives the growing importance of warehousing, Figure 5 shows the evolution of average labor productivity and average labor costs in OECD STAN subsector “Warehousing and support activities for transportation in the five largest EU Member States in our sample (Germany, France, Italy, Spain, The Netherlands) and the US. Average labor productivity is defined as value-added (in 2015 euros or dollars for all years) per employee. In the long-run, average labor productivity is expected to increase because of technological progress. Real average labor costs (in 2015 euros or dollars per unit output) are informative about the extent to which this productivity gain from technological progress is shared with workers through a higher average real wage. If

average labor productivity per worker increases relative to the average real wage, the labor share (i.e., the share of value-added that goes to workers) decreases. Alternatively, if average labor productivity per worker decreases relative to the average real wage, the labor share increases.


The solid black lines in Figure 5 plot the evolution of value-added per employee in warehousing, normalizing values to an index with 1995=100. For example, panel (a) for Germany shows that average labor productivity (i.e. value-added per employee) in warehousing was relatively low up to 2000, then increased rapidly to peak in 2006, during the 2000-2008 economic boom, fell during the recession of 2008-2010, and remained relatively stable thereafter. In the long-run, average labor productivity in warehousing increased by a substantial 38 percent between 1991 and 2019. The solid light blue line shows the evolution of average labor productivity for the entire economy. Eyeballing both the black and light blue solid lines suggests that warehousing is more pro-cyclical in the short-run and has stronger labor productivity growth in the long-run compared to the total German economy.

The dashed lines in Figure 5 plot the evolution of labor costs per employee (adjusted for productivity) in warehousing (the dashed orange line) and the total economy (the dashed green line). For example, panel (a) for Germany shows that the average real wage in warehousing increased 28 percent from 1995 to 2019. Before 2000, the average real wage in warehousing grew in line with economy-wide changes. After 2000, average real wage growth in warehousing was faster in the early 2000s, negative from 2005 to 2015 while average wages in the rest of the economy continued to grow, and again faster after 2015. In the long-run, average real wage growth in warehousing outpaced economy-wide average real wage growth, again pointing to the growing importance of warehousing in the German economy.


Figure 5. Evolution of average labor productivity and average labor costs in warehousing in several EU Member States and the US

Index (1995=100)


(a) Germany




(b) France




(c) Italy




(d) Spain



(e) Netherlands



(f) United States



Source: [OECD STAN database](#) for industrial analysis, 2020 ed.

Turning to the group of all five EU Member States, panels (a) to (e) of Figure 5 show some differences as well as some similarities between countries over the last three decades. The following three stylized facts summarize a comparison of all five European countries:

- 1) Average labor productivity in warehousing increased in Germany, France, and the Netherlands but held relatively constant in Spain and declined in Italy. In Germany, France, and the Netherlands,

average labor productivity in warehousing increased by more than the increase in economy-wide average labor productivity. The fast productivity growth of warehousing in some European countries is indicative of the development and successful adoption of new technologies in warehousing over time.

- 2) Average real wage changes in warehousing are more ambiguous across countries. The average real wage grew most in Germany (exceeding economy-wide growth), followed by the Netherlands. However, the average real wage stayed constant in the long run in France and decreased in Italy (but by less than productivity) as well as in Spain (but only in the late 1990s).
- 3) In all European countries except Italy, average labor productivity has grown faster than the average real wage in warehousing. This decoupling of productivity and wage growth is more outspoken for warehousing than for the entire economy. What this suggests is that only part or even none of the productivity gains in warehousing are shared with workers, resulting in a falling labor share. [Stansbury and Summers \(2020\)](#) examine several reasons why this could be the case. One reason for a decline in the labor share could be that warehouse technologies automate labor more than it augments what workers do. Another reason could be that the bargaining power of workers in warehousing has decreased over time, resulting in less rent sharing of the productivity gains from technological progress with workers. This could not only be driven by a decline in the unionization rates among warehouse workers, but also by the emergence of fissured labor contracts made possible by changes in labor law and labor market regulations.

Finally, panel (f) of Figure 5 shows the evolution of average labor productivity and the average real wage for warehousing and for the total economy in the US. In line with strong economy-wide productivity growth, average labor productivity in warehousing also increased 41 percent between 1995 and 2019. In contrast to the economy-wide decoupling of productivity and wage growth and fall in the labor share since 2000, average labor productivity and the average real wage in warehousing have shown a similar long-run trend.

Algorithmic management in warehousing

Productivity growth in warehousing could be driven by technological and/or organizational innovations. [Delfanti \(2019\) and Gent \(2018\) describe how Amazon distribution centers](#) are organized around four core processes: receive, store, pick, and pack. Receiving and storing are part of the “inbound” process, while picking and packing constitute the “outbound process.” Workers at receive stations unpack incoming pallets of commodities and identify each commodity by a unique barcode. Workers then store items in the pick area. The pick area usually is a large multi-floor area dense with thousands of shelves. When goods must be retrieved, workers walk through the warehouse to pick goods and carry them to packing stations. In the pack area, workers receive, package and label orders, which are then sent to shipping. Throughout these processes, goods are sometimes moved in plastic warehouse boxes (e.g., the yellow boxes in Amazon warehouses).

Delfanti (2019) writes that at the core of receiving, storing, picking and packing are product barcodes and various types of barcode scanners. These barcodes and scanners are not only used to keep inventory, but also to assist workers and collect information about workers' productivity. Workers pick up a portable barcode scanner at the beginning of their shift and use it to scan the barcode on their badge, thus logging into the system. From that moment on, the scanner mediates between workers and management, assigning tasks, communicating orders, and monitoring work. This is similar to forms of algorithmic management in the "gig economy," where the barcode scanner is replaced by the phone apps commonly used to collect and use data. Most often, these corporate algorithms are inherently opaque and access to them is prohibited by industrial secrecy and nondisclosure agreements. Therefore, auditing these algorithms is difficult.

Based on interviews with Amazon warehouse workers, [Delfanti \(2021\)](#) further describes the tasks of pickers in Amazon's MXP5 warehouse in Piacenza, a small town in northern Italy. A picker walks among the shelves pulling a cart which carries a box she needs to fill up with her batch of items to pick. Once an item has been picked, the picker uses her barcode scanner to scan the item's barcode. The barcode scanner records, approves, and communicates to the picker the next item she is to pick. It also communicates the position on the shelves and the time the picker has to complete the task—often a minute or so. One effect of this form of algorithmic management is that it augments what workers do: no individual human being can efficiently navigate an area of several thousand shelves to pick a list of wanted goods without the aid of an algorithm. Moreover, algorithmic management further increases labor productivity by requiring pickers to keep a fast "Amazon pace" (i.e., one cannot run but must walk as fast as possible).

However, the increased efficiency of warehouse operations due to data driven technologies also comes at a cost for warehouse workers. As the data that they generate is managed by algorithms and managers (who are often not working in the warehouse themselves but in some faraway global headquarter), warehouse workers lose agency over the tasks they do. Algorithms can dictate the pace and contents of work because they monopolize the knowledge about warehouse inventory and processing, while workers can only guess what data are being extracted and what analytics are being used to organize and surveil their activities. To illustrate this for the case of Amazon, [Delfanti \(2021\)](#) writes:

Technology dictates the pace of work at Amazon. It is used to increase workers' productivity, standardize tasks, facilitate worker turnover, and ultimately gain control over the workforce. Workers are acutely aware of the uneven nature of their relationship with machinery, and at the same time, they know the warehouse needs their living labor. As I was told by a manager, "technology codifies, understands, and manages. But the real machine is the human."

- *Delfanti (2021)*

Working conditions in warehousing

On the one hand, warehouses need workers because receiving, storing, picking, and packing still relies on workers' dexterity. Only humans possess the flexibility and speed to efficiently store items that are different in shape, weight, volume, color, and so on. On the other hand, workers need algorithms because what workers do not possess, or quickly surrender to algorithms, is their knowledge of the position of the commodities they have stored. In sum, workers and algorithmic management must co-exist in today's warehouses.

However, [Wood \(2021\)](#) argues that the balance of power between workers and algorithms may be increasingly tilted in favor of algorithms. As Figure 5 showed for all European countries in our sample except Italy, the productivity gains in warehousing are increasingly going to shareholders, in terms of dividends and stock options, at the expense of warehouse workers, in terms of their average compensation. This could be due to an increased use of algorithmic management combined with a decline in workers' bargaining power. For example, [Delfanti \(2021\)](#) argues that algorithmic management implies that minimal training for warehouse workers is needed to do their jobs. Usually, it only takes hours to train new associates to work as pickers. This enables warehouses to sustain high turnover rates while having access to a productive but also flexible workforce needed if there are sudden spikes in sales. But minimizing the costs of training and worker turnover is not enough: management also needs to ensure a smooth and productive relation between workers and algorithms on the warehouse floor. Therefore, workplace discipline and worker self-discipline are imposed by monitoring workers through the data they generate. To further cement worker effort, a workplace culture is created through slogans such as Amazon's "Work Hard. Have Fun. Make History."

Conclusions

During the past decades, the importance of warehousing has increased as lean manufacturing, lean retailing, and global-value chains spread across the economy. Today, employment in warehouses alone accounts for a sizeable 1% to 2% percent of total employment in advanced economies. Moreover, the emergence of digital technologies since the 1980s and, more recently, AI have transformed the nature of warehousing. Indicative of the successful adoption of new digital technologies in warehouses is an increase in average labor productivity and the average real wage.

However, growth in average labor productivity has outpaced average real wage growth, resulting in a fall in the labor share in warehousing. In modern fulfillment centers this could result from the use of algorithmic management together with a less strong representation of employees at the workplace. Despite strong growth in employment and wages in warehousing, these changing working conditions in warehouses pose an important challenge for workers. With advances in AI, the future of warehousing could be converging to algorithmic management systems that are fully

independent of workers. Fully automated warehouses that operate without the use of human labor are popularly known as “dark warehouses”. Dark warehouses imply that all the productivity gains in warehousing due to algorithmic management have entirely shifted away from warehouse workers (who are all displaced such that the labor share has fallen to zero). Although it is uncertain that most warehouses will become dark warehouses in the future, technologies currently developed for warehousing are in part geared towards automation instead of augmentation of human labor.

Part V: Conclusions

The use of AI undoubtedly presents many opportunities to positively transform the economy. The last decade has seen incredible advances in natural-language processing and computer vision, enabling new applications of AI to tasks previously thought to be firmly in the domain of humans. Firms are rapidly adopting AI around the world for its ability to scale and lower costs, to absorb and process enormous amounts of data, and to help make better decisions, often assisted by humans. And all this process of transition is likely to create new jobs that never would have existed without AI.

At the same time, AI poses several challenges. Huge swaths of the workforce are likely to be exposed to AI, in the sense that AI can now address nonroutine tasks, including tasks in high-skill jobs that until now had never been threatened by any kind of automation. The primary risk of AI to the workforce is in the general disruption it is likely to cause to workers, whether they find that their jobs are newly automated or that their job design has fundamentally changed. The additional risk of AI is that it may lead firms—unintentionally or not—to violate existing laws about bias, fraud, or antitrust, exposing themselves to legal or financial risk, and inflicting economic harm on workers and consumers. Given the black box nature of these systems, detecting, and addressing these violations is far from a simple task. This presents governments with a clear agenda on how to guide AI development in a positive direction.

- a) Investing in training and job transition services so that those employees most disrupted by AI can transition effectively to new positions where their skills and experience are most applicable.**

The introduction of AI across all areas of firms is likely to lead to large disruptions to the workforce. As seen in the case study on HR in part IV, AI can be a useful tool for helping workers find new opportunities with the same company by matching skills with job openings. This use of AI could help soften the disruption for some workers. However, there is likely to be a need for large investments in training, either to develop the new skills required for existing jobs that are being redesigned due to AI or for new jobs where there is growing demand for workers.

Long-term trends in employment complicate the idea of retraining workers. The increased prevalence of shorter contract durations lowers incentives for firms to invest in worker training, leading to underinvestment in skill acquisition. Policies that promote or subsidize intermediaries that share the costs and benefits of training are needed to reduce skill gaps, especially for workers at risk of automation. For example, temporary help agencies that provide training and match workers to employers for a fee avoid the flight risk that makes employers reluctant to invest in-house training, and can also split the cost of training across multiple employers or between

employers and workers. Employers pay a premium to such agencies for access to already trained workers, and workers receive part of their increased productivity in higher wages.

Intermediaries that invest in workers' skills to reduce skill gaps can be public, private, or hybrid. Examples include Public Employment Services, which offer training; outplacement offices, funded by companies that serve laid-off workers, which assist displaced workers in finding new jobs; and temporary help agencies specialized in training and finding jobs for workers who are otherwise unlikely to participate in the labor market. [Katz et al. \(2020\)](#) show that such policies could be particularly effective in increasing earnings and job mobility for trained workers.

Another set of programs that have been shown to increase wages focus on providing digital skills to workers. Companies report that the lack of staff with adequate digital skills is an obstacle to investment. Therefore, digital education action plans should ensure that more workers have basic digital skills. Digital skills not only include knowledge about computer science, technology, engineering, or mathematics (STEM), but also contain other skills that can complement new technologies. These other skills include communication and social skills that remain important competencies for workers even in workplaces that adopt AI systems.

b) Encouragement of development and adoption of AI that is beneficial for labor markets.

Firms, given their goal of maximizing profits, are most likely to conduct AI research and deploy AI systems that will directly benefit their bottom line. As a result, the development and adoption of AI is likely to diverge from what would be optimal for labor markets, i.e. workers' wages and employment. Three concerns emerge from our study:

1. Investing in the development of AI that augments workers: The most straightforward concern many workers face when it comes to AI is that of automation. [Acemoglu et al. \(2022\)](#) document that 54 percent of AI adopters do so to automate existing processes and a recurring theme raised by firms in the hiring space was that automation of certain elements of the HR profession was a key part of their business model and was being sought by their clients. Also, the case study of warehousing showed that algorithmic management of distribution centers is geared toward process and workforce automation. The use of AI to automate existing processes has important adverse effects on some workers, either through the full automation of their job or by substantially shifting the skill set required to perform the role. [Acemoglu \(2021\)](#) claims that while investment in AI and other technologies can lead to economic growth, firm incentives to reduce costs, increase shareholder value, and increase profits can direct firms away from a socially optimal split between automation and augmentation of worker tasks. One way to correct for private incentive-driven deviations from the optimal development of AI technology

is to utilize public funds to encourage and stimulate AI research that augments instead of automates work.

Further, in contrast to private efforts to develop AI, publicly funded academic research can focus on a wider array of AI methods and topics. This encompasses everything from an expanded set of AI methods, exploring the impact of AI on workers' wages and employment, the effects of algorithms on anti-competitive behavior in markets, the development of AI ethics, as well as how AI can exacerbate existing biases that are present in society, resulting in discrimination along racial, gender, and economic lines.

There is an ongoing effort along these lines to involve the public sector in AI research in both the US and the European Union. In the US, the National Security Commission on Artificial Intelligence, in its [final report](#), urged Congress to double federal R&D spending on AI each year, until it reached \$32 billion in 2026. The Biden Administration, in its [fiscal 2023 budget request](#), proposed increasing the federal R&D budget to more than \$204 billion, a 28 percent increase from 2021 enacted levels. Part of this funding would support new and existing National Artificial Intelligence Research Institutes. These institutes bring federal, state, and local agencies together with the private sector, nonprofits, and academia to tackle AI research. In Europe, [several funds support AI research](#), including Horizon Europe, the European Research Council, the European Innovation Council, and European Partnerships. An important focus in these ongoing efforts should be whether AI augments workers in their current jobs and whether AI creates new jobs for workers.

2. Public procurement of AI that augments workers: Public bodies can also direct the development of AI that complements workers through their procurement of AI systems. This public procurement can be aided by access to public data for AI developers. The availability of data can shape both the level and direction of innovative activity. This is explored by [Beraja et al. \(2022\)](#), who show that Chinese firms with access to data-rich government contracts develop substantially more commercial AI software.
3. Incentivizing private firms to adopt AI that augments workers: Beyond these ongoing efforts to fund AI research, governments have other mechanisms at their disposal to incentivize private firms to invest responsibly in AI. While public research efforts can be used to consciously prioritize AI research that improves worker productivity and encourages a diversity of techniques, governments should also be aware of the set of incentives faced by firms. These include firm business models that promote cutting costs, economic distortions in the tax and regulatory space that increase the cost to firms of using labor relative to capital, and even the “aspirations of researchers” at private firms who are excited and motivated to develop branches of AI that are more suited to automation ([Acemoglu 2021](#)). All these channels might push the society towards an

undesirable equilibrium in terms of the balance of automation and augmentation AI technologies.

c) Investing in the capacity of regulatory agencies to ensure that AI systems are transparent and fair for workers.

The case studies in part IV showed that AI systems can be made more transparent and fairer for workers in hiring processes and in algorithmic management of workplaces:

1. Algorithmic hiring: As discussed in part IV of this report, well-intentioned algorithms may still exhibit biases due to unforeseen implementation issues. And, as noted in the case study of AI and hiring, instances of bias are exacerbated when algorithms are being employed at nearly every step of a multistage process. And bias is not the only concern in this space. The black box nature of these tools means that there is a risk of fraud, where firms market their product as providing a service but their customers have no mechanism to determine the accuracy of their claims. Additionally, there is [evidence](#) that AI algorithms can learn to effectively collude with each other when setting prices. Existing discrimination, fraud, and antitrust rules and enforcement practices may be insufficient to counteract AI-created fraud and bias.
2. Algorithmic management of workplaces: The case study of warehousing also illustrates that workplace algorithms are inherently opaque, are shrouded in industrial secrecy, and are protected by nondisclosure agreements with workers. Even warehouse workers themselves are not always made aware of the software that manages them. This reflects workers' own relation with corporate technology: workers' relations with AI are based on information asymmetries, given that workers can only guess the procedures of data extraction and analytics that organize and surveil their activities.

Firms that utilize AI are not freed of the responsibility of abiding by antifraud, antitrust, and antidiscrimination laws, as well as workplace safety and health regulations. It should be a principal goal of policymakers to make sure that government institutions are well-equipped to investigate and enforce these laws when necessary. Doing so is not a straightforward process. A recent Brookings Institution [report](#) highlighted several necessary steps: creating robust standards for algorithmic audits, assuring that [regulatory agencies](#) have the access to firms when they need to perform an audit, building technical expertise within agencies, and [revising and crafting policies](#) that are aware of the challenges of overseeing algorithm-driven workplaces. The goal is to create the appropriate incentives so that firms develop fairer algorithms that abide by national laws. As noted in the hiring case study, well-designed algorithms have the potential to actual reduce instances of bias, and firms have [voiced](#) a desire to use algorithms to address instances of

discrimination. However, without the proper oversight and regulation, this potential for positive transformative change is unlikely to be realized.

Governments are already moving toward more effective regulation of the impact of AI. In October 2022, Spain launched [a pilot regulatory sandbox on AI](#). This sandbox is a way to connect policymakers with AI developers and adopters. It is expected to generate easy-to-follow best practice guidelines for companies, including small and medium-size enterprises and start-ups, to stimulate the development of and reduce barriers to adopt AI, in compliance with the future [European Commission Artificial Intelligence Act](#). Further, the US has announced an initiative to create a “[Bill of Rights](#)” for AI covering many areas, such as consumer protections and equity of opportunity in employment, education, housing and finance, and health care.

References

- Abid A., M. Farooqi, and J. Zhou. 2021. “Large Language Models Associate Muslims with Violence.” *Nature Machine Intelligence* 3: 461-463.
<https://www.nature.com/articles/s42256-021-00359-2>
- Acemoglu, D. 2021. “Harms of AI.” NBER Working Paper 29247. Cambridge, MA: National Bureau of Economic Research. <https://www.nber.org/papers/w29247>
- Acemoglu D., and D. Autor. 2011. “Skills, Tasks and Technologies: Implications for Employment and Earnings.” *Handbook of Labor Economics*, 4: 1043-1171.
[https://doi.org/10.1016/S0169-7218\(11\)02410-5](https://doi.org/10.1016/S0169-7218(11)02410-5)
- Acemoglu, D., D. Autor, J. Hazell, and P. Restrepo. 2022. “Artificial Intelligence and Jobs: Evidence from Online Vacancies.” *Journal of Labor Economics*, 40.
<https://doi.org/10.1086/718327>
- Acemoglu D., G. Anderson, D. Beede et al. 2022. “Automation and the Workforce: A Firm-Level View From the 2019 Annual Business Survey.” Paper presented at the NBER/CRIW conference on Technology, Productivity and Economic Growth, Washington DC, March 2022. <http://pascual.scripts.mit.edu/research/abs/>
- Aiken, E., S. Bellue, D. Karlan, et al. 2022. “Machine Learning and Mobile Phone Data Can Improve the Targeting of Humanitarian Assistance.” NBER Working Paper 29070. Cambridge, MA: National Bureau of Economic Research.
https://www.nber.org/system/files/working_papers/w29070/w29070.pdf
- Albernathy, F., J. Dunlop, J. Hammond, D. Weil. 2000. “Control Your Inventory in a World of Lean Retailing.” *Harvard Business Review*, November-December, 2000.
<https://hbr.org/2000/11/control-your-inventory-in-a-world-of-lean-retailing>
- Center for Democracy and Technology. 2020. “Algorithm-driven Hiring Tools.”
<https://cdt.org/wp-content/uploads/2020/12/Full-Text-Algorithm-driven-Hiring-Tools-Innovative-Recruitment-or-Expedited-Disability-Discrimination.pdf>
- AlphaFold Protein Structure Database. “AlphaFold Protein Structure Database.”
<https://alphafold.ebi.ac.uk/>.
- DeepMind. “AlphaGo.” <https://www.deepmind.com/research/highlighted-research/alphago>

- Ameri, M., L. Schur, M. Adya et al. 2017. “The Disability Employment Puzzle: A Field Experiment on Employer Hiring Behavior.” *SAGE Journals ILR Review*, 71, no.2. <https://doi.org/10.1177/0019793917717474>
- Autor, D. 2022. “The Labor Market Impact of Technological Change: From Unbridled Enthusiasm to Qualified Optimism to Vast Uncertainty.” NBER Working Paper 30074. Cambridge, MA: National Bureau of Economic Research. <https://www.nber.org/papers/w30074>
- Autor, D., C. Chin, A. Salomons, and B. Seegmiller. 2022. “New Frontiers: The Origin and Content of New Work 1940-2018.” NBER Working Paper 30389. Cambridge, MA: National Bureau of Economic Research. <https://www.nber.org/papers/w30389>
- Bennet, M., C. Booker, S. Brown et al. 2020. *Letter to Janet Dhillon*. United States Senate. https://www.bennet.senate.gov/public/_cache/files/0/a/0a439d4b-e373-4451-84ed-ba333ce6d1dd/672D2E4304D63A04CC3465C3C8BF1D21.letter-to-chair-dhillon.pdf
- Beraja, M., A. Kao, D. Yang, N. Yuchtman. 2022. “AI-tocracy.” Working paper, MIT Economics. https://economics.mit.edu/sites/default/files/2022-09/aitocracy_20220701.pdf
- Beraja, M., D. Yang, and N. Yuchtman. 2022. “Data-intensive Innovation and the State: Evidence from AI Firms in China.” Working paper, MIT Economics. https://economics.mit.edu/sites/default/files/publications/dataproduction_draft_20220110.pdf
- Bessen, J., M. Goos, A. Salomons, W. Van den Berge. 2019. “Automation Reaction – What Happens to Workers at Firms that Automate?” Boston University School of Law, Law and Economics Research Paper. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3328877
- Brown, S. 2021. “Machine Learning, Explained.” MIT Sloan. <https://mitsloan.mit.edu/ideas-made-to-matter/machine-learning-explained>
- Brynjolfsson, E., A. Collis, E. Diewert, F. Eggers, and K. Fox. 2019. “GDP-B: Accounting for the Value of New and Free Goods in the Digital Economy.” NBER Working Paper 25695. Cambridge, MA: National Bureau of Economic Research. <https://www.nber.org/papers/w25695>

- Brynjolfsson, E., T. Mitchell, and D. Rock. 2018. “What Can Machines Learn, and What Does it Mean for Occupations and the Economy?” *AEA Papers and Proceedings*, 108: 43-47.
<https://doi.org/10.1257/pandp.20181019>
- Calvano, E., G. Calzolari, V. Denicolo et al. 2020. “Protecting Consumers from Collusive Prices due to AI.” *Science*, 370, no. 6520: 1040-1042.
<https://www.science.org/doi/full/10.1126/science.abe3796>
- Centers for Medicare and Medicaid Services. 2021. “Artificial Intelligence (AI) Health Outcomes Challenge.” Centers for Medicare and Medicaid Services.
<https://innovation.cms.gov/innovation-models/artificial-intelligence-health-outcomes-challenge>
- Centers for Medicare and Medicaid Services. 2021. “CMS Selects Winner and Runner-Up in Artificial Intelligence Health Outcomes Challenge.” *Newsroom*, April 30, 2021.
<https://www.cms.gov/newsroom/press-releases/cms-selects-winner-and-runner-artificial-intelligence-health-outcomes-challenge>
- Cockburn, I., R. Henderson, and S. Stern. 2019. “The Impact of Artificial Intelligence on Innovation: An Exploratory Analysis,” In *The Economics of Artificial Intelligence: An Agenda*, edited by A. Agrawal, J. Gans, and A. Goldfarb, 115-146. Chicago: University of Chicago Press. <https://www.nber.org/books-and-chapters/economics-artificial-intelligence-agenda/impact-artificial-intelligence-innovation-exploratory-analysis>
- OpenAI. “Dall·E 2.” <https://openai.com/dall-e-2/>
- Dean, J. 2019. “The Deep Learning Revolution and its Implications for Computer Architecture and Chip Design.” Working Paper, Google Research.
<https://arxiv.org/ftp/arxiv/papers/1911/1911.05289.pdf>
- Delfanti, A. 2019. “Machinic Dispossession and Augmented Despotism: Digital Work in an Amazon Warehouse.” *SAGE Journals*, 23, no.1.
<https://doi.org/10.1177/1461444819891613>
- Delfanti, A. 2021. *The Warehouse: Workers and Robots at Amazon*. London: Pluto Press
<https://www.jstor.org/stable/j.ctv2114fnm>
- Engler, A. 2021. “Auditing Employment Algorithms for Discrimination.” The Brookings Institution, Washington, DC. <https://www.brookings.edu/research/auditing-employment-algorithms-for-discrimination/>

- European Commission. 2020. “European Enterprise Survey on the Use of Technologies Based on Artificial Intelligence: Final Report.” Directorate-General for Communications Networks, Content. <https://data.europa.eu/doi/10.2759/759368>
- European Commission, 2021, “EU-US Trade and Technology Council Inaugural Joint Statement”,
https://ec.europa.eu/commission/presscorner/detail/en/STATEMENT_21_4951
- European Commission. 2022. “First Regulatory Sandbox on Artificial Intelligence Presented.”
<https://digital-strategy.ec.europa.eu/en/news/first-regulatory-sandbox-artificial-intelligence-presented>
- European Commission. 2022. “A European Approach to Artificial Intelligence.” The European Union. <https://digital-strategy.ec.europa.eu/en/policies/european-approach-artificial-intelligence>
- European Commission. *No date*. “Research and Innovation: Artificial Intelligence (AI).”
https://research-and-innovation.ec.europa.eu/research-area/industry/key-enabling-technologies/artificial-intelligence-ai_en
- Eurostat. 2017. “Glossary: Community Survey on ICT Usage in Enterprises.”
https://ec.europa.eu/eurostat/statisticsexplained/index.php?title=Glossary:Community_survey_on_ict_usage_in_enterprises
- Eurostat. 2021. “Artificial intelligence in EU enterprises.” *Eurostat News*, April 13, 2021.
<https://ec.europa.eu/eurostat/web/products-eurostat-news/-/ddn-20210413-1>
- Eurostat. 2022. “Digital Economy and Society Statistics – Enterprises.” *Eurostat Statistics Explained*, January 20, 2022. https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Digital_economy_and_society_statistics_-_enterprises
- Feigenbaum, J. and D. Gross. 2020. “Answering the Call of Automation: Labor Market Adjusted to the Mechanization of Telephone Operation.” NBER Working Paper 28061. Cambridge, MA: National Bureau of Economic Research.
<https://www.nber.org/papers/w28061>
- Felten, E., M. Raj, R. Seamans. 2019. “The Occupational Impact of Artificial Intelligence: Labor, Skills, and Polarization.” Working Paper, NYU Stern School of Business.
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3368605

- Forbes Insights. 2019. “How Pandora Knows What you Want to Hear Next.” *Forbes*, October 1, 2019. <https://www.forbes.com/sites/insights-teradata/2019/10/01/how-pandora-knows-what-you-want-to-hear-next/?sh=1e6d5f613902>
- Fuller, J. and W. Kerr. 2022. “The Great Resignation Didn’t Start with the Pandemic.” *Harvard Business Review*, March 23, 2022. <https://hbr.org/2022/03/the-great-resignation-didnt-start-with-the-pandemic>
- Data & Trust Alliance. “Data & Trust Alliance.” <https://dataandtrustalliance.org/>
- Gaynor, M. 2020. “Automation and AI Sound Similar, but May Have Vastly Different Impacts on the Future of Work.” The Brookings Institution Blog, January 30, 2020. <https://www.brookings.edu/blog/the-avenue/2020/01/29/automation-and-artificial-intelligence-sound-similar-but-may-have-vastly-different-impacts-on-the-future-of-work/>
- Gender Shades “Gender Shades: Results.” <http://gendershades.org/overview.html>
- Gent, C. 2018. “The Politics of Algorithmic Management Class: Composition and everyday struggle in distribution work.” Warwick, *University of Warwick*. <http://wrap.warwick.ac.uk/132957/>
- Genz, S., T. Gregory, M. Janser, F. Lehmer, and B. Matthews. 2021. “How do Workers Adjust When Firms Adopt New Technologies?” Working Paper, ZEW – Centre for European Economic Research Discussion Paper no. 21-073. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3949800
- Gershgorn, D. 2022. “Companies are on the Hook if Their Hiring Algorithms are Biased.” *Quartz*, July 20, 2022. <https://qz.com/1427621/companies-are-on-the-hook-if-their-hiring-algorithms-are-biased/>
- Goger, A., A. Parco, E. Vegas. 2022. “Learning and Working in the Digital Age.” The Brookings Institution, Washington DC. https://www.brookings.edu/wp-content/uploads/2022/05/Learning-and-working-in-the-digital-age_FINAL.pdf
- Goos M., A. Manning, and A. Salomons. 2014. “Explaining Job Polarization: Routine-Biased Technological Change and Offshoring.” *American Economic Review*, 104 no. 8: 2509-2526.
- Hardesty, L. 2017. “Explained: Neural Networks.” *MIT News Office*, April 14, 2017. <https://news.mit.edu/2017/explained-neural-networks-deep-learning-0414>

- Hoffman, B. 2018. “Competition and Consumer Protection Implications of Algorithms, Artificial Intelligence and Predictive Analytics.” Remarks at Competition and Consumer Protection in the 21st Century.
https://www.ftc.gov/system/files/documents/public_statements/1431041/hoffman_-_ai_intro_speech_11-14-18.pdf
- Hoffmann, M., L. Nurski. 2021. “What is Holding Back Artificial Intelligence Adoption in Europe?” *Bruegel*, Nov 30, 2021. <https://www.bruegel.org/policy-brief/what-holding-back-artificial-intelligence-adoption-europe>
- Digital Credentials Consortium. “Digital Credentials Consortium.”
<https://digitalcredentials.mit.edu/>
- Hutchinson, B., V. Prabhakaran, E. Denton, K. Webster, Y. Zhong, and S. Denuyl. 2019. “Unintended Machine Learning Biases as Social Barriers for Persons with Disabilities.” *Special Interest Group on Accessible Computing Newsletter*, October 2019.
<http://sigaccess.org/newsletter/2019-10/hutchinson.html>
- ImageNet. ImageNet Large Scale Visual Recognition Challenge 2010.” <https://www.image-net.org/challenges/LSVRC/2010/index.php>
- ImageNet “ImageNet Large Scale Visual Recognition Challenge 2017.”. <https://image-net.org/challenges/LSVRC/2017/results>
- Indeed. *No date*. “Pros and Cons of Recruiting Chatbots.”
<https://www.indeed.com/hire/c/info/pros-and-cons-of-recruiting-chatbots#3>
- IRS. 2022. “Assistance for Eligible Taxpayers in Setting up or Modifying Payment Plans Now Available.” *Internal Revenue Service Newsroom*, June 17.
<https://www.irs.gov/newsroom/irs-expands-voice-bot-options-for-faster-service-less-wait-time>
- Jumper, J., R. Evans, A. Pritzel, et al. 2021. “Highly Accurate Protein Structure Prediction with AlphaFold.” *Nature* 596: 583-589. <https://www.nature.com/articles/s41586-021-03819-2>
- Katz, L., J. Roth, R. Hendra, and K. Schaberg. 2020. “Why do Sectoral Employment Programs Work? Lessons from WorkAdvance.” NBER Working Paper 28248. Cambridge, MA: National Bureau of Economic Research. <https://www.nber.org/papers/w28248>

- Kerry, C. 2020. “Protecting Privacy in an AI-driven World.” The Brookings Institution, Washington DC. <https://www.brookings.edu/research/protecting-privacy-in-an-ai-driven-world/>
- Kleinberg, J., H. Lakkaraju, J. Leskovec, et al. 2018. “Human Decisions and Machine Predictions.” *The Quarterly Journal of Economics* 133, no. 1: 237-293. <https://doi.org/10.1093/qje/qjx032>
- Koenecke, A., A. Nam, E. Lake, J. Nudell et al. 2020. “Racial disparities in automated speech recognition.” *PNAS* 117, no. 14: 7684-7689. <https://doi.org/10.1073/pnas.1915768117>
- Krayewski, K. 2022. “How NLP Chatbots Work.” *Ultimate AI Blog*, April 13, 2022. <https://www.ultimate.ai/blog/ai-automation/how-nlp-text-based-chatbots-work>
- Lambrech, A., C. Tucker. 2019. “Algorithmic Bias? An Empirical Study of Apparent Gender-Based Discrimination in the Display of STEM Career Ads.” *Management Science*. <https://doi.org/10.1287/mnsc.2018.3093>
- Martinez, E. and L. Kirchner. 2021. “The Secret Bias Hidden in Mortgage-Approval Algorithms.” *The Markup*, August 25, 2021. <https://themarkup.org/denied/2021/08/25/the-secret-bias-hidden-in-mortgage-approval-algorithms>
- Meta. 2020. “Good Questions, Real Answers: How Does Facebook Use Machine Learning to Deliver Ads?” *Announcements*, June 11, 2020. <https://www.facebook.com/business/news/good-questions-real-answers-how-does-facebook-use-machine-learning-to-deliver-ads>
- Mitchell, M. 2021. “Why AI is Harder Than We Think.” Working Paper, Santa Fe Institute. <https://arxiv.org/pdf/2104.12871.pdf>
- Mokyr, J. 2005. “Long-term Economic Growth and the History of Technology”. *Handbook of Economic Growth*, Volume 1, Part B, 2005, Pages 1113-1180. <https://www.sciencedirect.com/science/article/abs/pii/S1574068405010178>
- National Artificial Intelligence Initiative Office. *No Date*. “Strategic Pillars: Applications.” <https://www.ai.gov/strategic-pillars/applications/>
- The National Security Commission on Artificial Intelligence. “The National Security Commission on Artificial Intelligence.” <https://www.nsc.ai.gov/>

- Obermeyer Z., and B. Powers, C. Vogelli, and S. Mullainathan. 2019. “Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations.” *Science*, 366: 447-453.
<https://www.science.org/doi/10.1126/science.aax2342>
- Organization for Economic Cooperation and Development (OECD). 2019. “OECD AI Principles Overview.” <https://oecd.ai/en/ai-principles>
- Organization for Economic Cooperation and Development (OECD). “STAN STructural ANalysis Database.” OECD.
<https://www.oecd.org/sti/ind/stanstructuralanalysisdatabase.htm>
- Organization for Economic Cooperation and Development (OECD). 2022. “Trends and Data Overview.” <https://oecd.ai/en/trends-and-data>
- Olmstead, K. 2017. “Nearly half of Americans Use Digital Voice Assistants mostly on their smartphones.” *Pew Research Center Fact-Tank*, December 12, 2017.
<https://www.pewresearch.org/fact-tank/2017/12/12/nearly-half-of-americans-use-digital-voice-assistants-mostly-on-their-smartphones/>
- Olson, C., and K. Kemery. 2019. “Voice report: From answers to action: customer adoption of voice technology and digital assistants.” Microsoft. https://advertiseonbing-blob.azureedge.net/blob/bingads/media/insight/whitepapers/2019/04%20apr/voice-report/bingads_2019_voicereport.pdf
- OpenAI “Overview.”. <https://openai.com/api/>
- Pethokoukis, J. 2016. “What the Story of ATMs and Bank Tellers Reveals About the ‘Rise of the Robots’ and Jobs.” *AEI Ideas*, June 6, 2016. <https://www.aei.org/economics/what-atms-bank-tellers-rise-robots-and-jobs/>
- Quillian, L., D. Pager, O. Hexel, and A. Midtboen. 2017. “Meta-analysis of Field Experiments Shows No Change in Racial Discrimination in Hiring Over Time.” *Proceedings of the National Academy of Sciences*, 114 no. 41:10870-10875.
<https://doi.org/10.1073/pnas.1706255114>
- Raghavan, M., S. Barocas. 2019. “Challenges for Mitigating Bias in Algorithmic Hiring.” The Brookings Institution, Washington DC. <https://www.brookings.edu/research/challenges-for-mitigating-bias-in-algorithmic-hiring/>
- Ramesh, A., M. Pavlov, G. Goh, et al. 2021. “Zero-Shot Text-to-Image Generation.” Working Paper, Cornell University Arxiv. <https://arxiv.org/abs/2102.12092>

- Ramesh, A., M. Pavlov, G. Goh, and S. Gray. 2021 “Creating Images from Text.” *Open AI Blog*, January 25, 2021. <https://openai.com/blog/dall-e/>
- Russakovsky, O., J. Deng, H. Su, et al. 2015. “ImageNet Large Scale Visual Recognition Challenge.” *J Comput Vis* 115: 211-252. <https://doi.org/10.1007/s11263-015-0816-y>
- Silberg, J., J. Manyika. 2019. “Tackling Bias in Artificial Intelligence (and in Humans).” *McKinsey Global Institute*, June 6. <https://www.mckinsey.com/featured-insights/artificial-intelligence/tackling-bias-in-artificial-intelligence-and-in-humans>
- Silver, D., J. Schrittwieser, K. Simonyan, et al. 2017. “Mastering the Game of Go Without Human Knowledge.” *Nature* 550: 354-359. <https://doi.org/10.1038/nature24270>
- Stansbury, A., L. Summers. 2020. “The Declining Worker Power Hypothesis: An Explanation for the Recent Evolution of the American Economy.” The Brookings Institution, Washington DC. <https://www.brookings.edu/wp-content/uploads/2020/12/StansburySummers-Final-web.pdf>
- Textio. “Inclusive Talent Attraction.” <https://textio.com/products/recruiting>
- U.S. Bureau of the Census. 2019. “Annual Business Survey.” <https://www.census.gov/programs-surveys/abs/data/tables.html>
- U.S. Department of Transportation. 2022. “The Economic Costs of Freight Transportation.” https://ops.fhwa.dot.gov/freight/freight_analysis/freight_story/costs.htm
- Webb, M. 2020. “The Impact of Artificial Intelligence on the Labor Market.” Working Paper, Stanford University. https://www.michaelwebb.co/webb_ai.pdf
- Weil, D. 2019. “Preparing for the Future of Work Through Understanding the Present of Work: A Fissured Workplace Perspective.” *Testimony before the U.S. House of Representatives Committee on Education and Labor and Subcommittee on Workforce Protections and Subcommittee on Health, Employment, Labor, and Pensions*. <https://edlabor.house.gov/imo/media/doc/WeilTestimony102319.pdf>
- Weldon, S. and Liang, Y. 2021. “Translation API Advanced can Translate Business Documents Across 100+ Languages.” *Google Cloud Blog*, May 12, 2021. <https://cloud.google.com/blog/products/ai-machine-learning/google-cloud-ai-translation-services-works-for-documents>
- The White House. 2022. “Blueprint for an AI Bill of Rights.” The United States Government. <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>

- The White House. 2022. “Fiscal 2023 Budget Chapter 18: Research And Development”
https://www.whitehouse.gov/wp-content/uploads/2022/04/ap_18_research_fy2023.pdf
- The White House. 2021. “U.S. – EU Trade and Technology Council Inaugural Joint Statement.”
Briefing Room, September 29, 2021. <https://www.whitehouse.gov/briefing-room/statements-releases/2021/09/29/u-s-eu-trade-and-technology-council-inaugural-joint-statement/>
- Wood, A. 2021. “Algorithmic Management.” European Commission, JRC Technical Report.
<https://joint-research-centre.ec.europa.eu/system/files/2021-05/jrc124874.pdf>
- Zolas N., Z. Kroff, E. Brynjolfsson et al. 2020. “Advanced Technologies Adoption and Use by Firms: Evidence from the Annual Business Survey.” NBER Working Paper 28290.
Cambridge, MA: National Bureau of Economic Research.
<https://www.nber.org/papers/w28290>