


This PDF is available at <http://nap.nationalacademies.org/27644>



## Artificial Intelligence and the Future of Work (2025)

### DETAILS

200 pages | 8.5 x 11 | PAPERBACK

ISBN 978-0-309-71714-4 | DOI 10.17226/27644

### CONTRIBUTORS

Committee on Automation and the U.S. Workforce: An Update; Computer Science and Telecommunications Board; Division on Engineering and Physical Sciences; Board on Human-Systems Integration; Division of Behavioral and Social Sciences and Education; National Academies of Sciences, Engineering, and Medicine

### SUGGESTED CITATION


National Academies of Sciences, Engineering, and Medicine. 2025. *Artificial Intelligence and the Future of Work*. Washington, DC: The National Academies Press. <https://doi.org/10.17226/27644>.

BUY THIS BOOK

FIND RELATED TITLES

Visit the National Academies Press at [nap.edu](http://nap.edu) and login or register to get:

- Access to free PDF downloads of thousands of publications
- 10% off the price of print publications
- Email or social media notifications of new titles related to your interests
- Special offers and discounts



All downloadable National Academies titles are free to be used for personal and/or non-commercial academic use. Users may also freely post links to our titles on this website; non-commercial academic users are encouraged to link to the version on this website rather than distribute a downloaded PDF to ensure that all users are accessing the latest authoritative version of the work. All other uses require written permission. ([Request Permission](#))

This PDF is protected by copyright and owned by the National Academy of Sciences; unless otherwise indicated, the National Academy of Sciences retains copyright to all materials in this PDF with all rights reserved.

# Summary

Interest in how advances in artificial intelligence (AI) will affect workers has been growing in recent years, especially with the rapid increase in capabilities and adoption of large language model (LLM)-based chatbots and other generative AI.

This report, requested in Section 5105 of the 2021 National Defense Authorization Act,<sup>1</sup> builds on a 2017 National Academies of Sciences, Engineering, and Medicine study<sup>2</sup> that examined the impacts of information technology on the workforce. It reviews current knowledge about the workforce implications of AI and related technologies, identifies key open questions, and describes salient research opportunities and data needs. The key findings are as follows:

Finding 1: AI is a general-purpose technology<sup>3</sup> that has recently undergone significant rapid progress. Still, there is a great deal of uncertainty about its future course, suggesting that wide error bands and a range of contingencies should be considered.

Finding 2: AI systems today remain imperfect in multiple ways. For example, LLMs can “hallucinate” incorrect answers to questions, exhibit biased behavior, and fail to reason correctly to reach conclusions from given facts.

---

<sup>1</sup> Public Law (P.L.) 116-283.

<sup>2</sup> National Academies of Sciences, Engineering, and Medicine, 2017, *Information Technology and the U.S. Workforce: Where Are We and Where Do We Go from Here?* The National Academies Press, <https://doi.org/10.17226/24649>.

<sup>3</sup> General-purpose technologies like the steam engine and electricity have widespread applications and were thus key drivers of economic growth. As discussed in more detail later, AI is advancing exceptionally rapidly, reflecting several key technical breakthroughs.

Finding 3: Significant further advances in AI technology are highly likely, but experts do not agree on the exact details and timing of likely advances.

Finding 4: The substantial and ongoing improvements in AI's capabilities, combined with its broad applicability to a large fraction of the cognitive tasks in the economy and its ability to spur complementary innovations, offer the promise of significant improvements in productivity.

Finding 5: As was the case with earlier general-purpose technologies, achieving the full benefits of AI will likely require complementary investments in new skills and new organizational processes and structures.

Finding 6: The labor market consequences of widespread AI deployment will depend both on the rate at which AI's capabilities evolve and on demographic, social, institutional, and political forces that are not technologically determined.

Finding 7: AI can be used to improve worker outcomes or to displace workers. Too often an exclusive focus on worker displacement neglects two other potentially positive labor market consequences of AI—new forms of work that demand valuable new expertise and AI systems that work jointly with workers to enable them to use their expertise more effectively to accomplish a broader variety of valuable tasks, perhaps with less formal training.

Finding 8: History suggests that even if AI yields significantly higher worker productivity, the productivity gains might fall unevenly across the workforce and might not be reflected in broad-based wage growth.

Finding 9: AI will have significant implications for education at all levels, from primary education, through college, through continuing education of the workforce. It will drive the demand for education in response to shifting job requirements, and the supply of education as AI provides opportunities to deliver education in new ways. It may also shift what is taught to the next generation to prepare them to take full advantage of future AI tools and advances.

Finding 10: Better measurement of how and when AI advancements affect the workforce is needed. To help workers adapt to a changing world, improving the ability to observe and communicate these changes—such as the impact of LLMs on knowledge work and robotics on physical work—as they occur is crucial.

Finding 11: Responses to concerns that AI poses potentially serious risks in areas such as fairness, bias, privacy, safety, national security, and civil discourse will modulate the rate and extent of impact on the workforce. It will take deep technical knowledge and may require new institutional forms for governments to stay abreast of and address these issues, given the rapidly changing technology.

## ARTIFICIAL INTELLIGENCE

The year following ChatGPT's introduction in November 2022 marked a major inflection point for AI with the emergence of and widespread public access to generative AI, especially LLMs. These LLMs exhibit major new AI capabilities compared to earlier AI systems, including the ability to hold meaningful conversations about diverse topics in dozens of languages, summarize the key points discussed in large text documents, perform a variety of problem-solving tasks, and write computer programs.

This rapid rate of AI progress is likely to continue for some years, owing to expected large commercial and government investments to develop bigger and better models, the availability of increasingly diverse and ever larger data sets to train AI systems, progress in open-source efforts to develop more shareable and portable models, and a burst of effort by both start-ups and mature corporations to apply this technology to a wide range of applications. There are also forces that work to slow the rate of advance, such as the need to address shortcomings of this imperfect technology, the need for the technology to be socially acceptable and trusted by the public (e.g., to avoid implicit social biases or to avoid helping bad actors achieve harmful goals), potential government regulation, and decisions by companies to limit access to AI capabilities in light of these and other challenges, including privacy concerns.

Other areas of AI—including speech recognition; computer vision and other forms of computer perception; the application of machine learning to large, structured data sets; autonomous vehicles and robotics—are experiencing slower but continued technical progress. There is some possibility that future advances in LLMs and extensions such as multimodal foundation models trained on a combination of text, images, videos, voices, and sounds may spill over to accelerate progress in these areas as well, although this is not guaranteed.

There is great uncertainty regarding which specific AI capabilities will appear in the coming years and when. As a result, decision makers need to create policies that will be robust to a variety of possible future technology advances and timetables. Moreover, the capabilities developed and, more importantly, whether and how they are implemented will depend on society's collective choices.

## ARTIFICIAL INTELLIGENCE AND PRODUCTIVITY

Productivity growth—growth in the amount of output per unit input—is crucial for improving long-term living standards. Advances in general-purpose technologies are key to enhancing productivity. Although AI is a general-purpose technology—that is, one that is pervasive, improves over time, and spawns complementary innovations—offering substantial productivity improvements in specific tasks and with the potential for impact in every economic sector, its overall impact on aggregate productivity remains minimal, potentially reflecting its low adoption rate across the economy thus far. However, given AI’s wide applicability and the rate of its technical development, significant impacts on productivity are likely in the coming decade.

Improvements in productivity stem from finding more effective ways to use labor and other inputs. Technological progress, especially in the form of general-purpose technologies like AI, is the primary driver of long-term economic growth, influencing investment in physical capital and improving productivity. The productivity effects of AI, which can substitute and complement labor, depend on how AI affects different sectors and tasks.

Generative AI has already been found to increase productivity in specific applications including contact centers, software development, and writing. Of course, many other tasks are not suitable for AI, at least in its current form. Considering the set of tasks potentially affected and the factors on which productivity effects depend, AI has the potential to increase aggregate productivity growth substantially for the broader economy in the coming decade. Although it is notoriously difficult to predict the details of future impact, some estimates suggest as much as a doubling of the rate of growth in the U.S. gross domestic product (GDP) from about 1.4 percent currently to 3 percent. Moreover, generative AI’s potential to accelerate scientific discovery and innovation could further compound productivity gains. However, tremendous uncertainty remains about the exact magnitude and timing of any productivity gains, and productivity could be hurt in some cases by increased fraud, misinformation, or other dangers.

Even if the productivity gains are large, there is no guarantee that these benefits will be distributed equitably. Without institutional and policy changes, the benefits might not be shared widely, potentially leading to job losses and wage disparities, increased inequality, and adverse effects on job quality and worker satisfaction. As productivity is not the only measure of human well-being, it is important to consider how AI might affect other aspects of human well-being, such as social progress and happiness as well as societal risks associated with AI. The latter includes threats to privacy, the potential for discrimination and bias, risks to democracy and political stability, and national security concerns—all of which will require attention by business leaders and policy makers.

## ARTIFICIAL INTELLIGENCE AND THE WORKFORCE

Although there is widespread concern about the impacts of AI on jobs, at the time of this writing U.S. unemployment rates are very low compared to historical levels; apart from a spike in unemployment owing to the COVID-19 pandemic, they have been extremely low for the past several years. In addition, population and labor force growth rates in the United States and across the industrialized world are expected to decelerate. Against this backdrop of structurally strong demand for labor and structural headwinds impeding increases in labor supply, it is difficult to predict whether adoption of new AI will result in a decline in aggregate labor demand, manifesting in either fewer jobs (relative to working-age population) or, more likely, lower pay for existing work.<sup>4</sup> Additionally, adoption of recent AI advances in the workplace is still nascent and—despite recent improvements—measurement of AI’s impacts is still limited, precluding a definitive assessment of the current impacts of AI on the workforce.

What is easier to predict is that adoption of AI will alter the nature of jobs—the set of tasks that define the job, and the share of these tasks that will be done by AI or in collaboration between human workers and AI tools. As a result, the key question is how AI will alter the demand for different types of worker expertise. (“Expertise” here refers to a specific body of knowledge or competency required to accomplish a particular objective—for example, baking bread, taking vital signs, or coding an app.)

New technology can erode the value of existing types of expertise (e.g., tax preparation software erodes the value of tax expertise) or create opportunities for jobs that require new kinds of expertise (e.g., computers create new jobs that require software engineering expertise). The creation of demand for new expertise is a critical force that counterbalances the tendency of new technology to erode the value of old expertise.

The future impact of AI on the demand for expertise is uncertain, but three plausible scenarios emerge. First, AI could accelerate occupational polarization, automating more nonroutine tasks and increasing the demand for elite expertise while displacing middle-skill workers. Second, AI might advance to outcompete humans across nearly all domains, greatly reducing the value of human labor and creating significant income distribution challenges. However, the committee believes that this scenario is unlikely in the near future owing to the limitations of AI, demographic trends, and the potential for new forms of work to emerge.

---

<sup>4</sup> Because most people are dependent on labor income for their livelihoods, a fall in earnings does not unambiguously imply a fall in employment; it is possible that more workers will work more hours to meet economic needs. Thus, earnings rather than employment are likely to be a better summary measure of labor demand.

Third, a more speculative scenario envisions a future in which the demand for expertise borrows attributes from both elite and mass expertise, leading to a reinstatement of the value of mass expertise in new domains. The creation of demand for novel human expertise that possesses market value and augments labor demand is a central attribute of the process of technological change but often is underemphasized compared to the countervailing force of automation. In the two major technological transitions preceding the AI revolution (from the artisanal era to the industrial era, and from the industrial era to the computer era), important categories of previously valuable human expertise were stranded—that is, made economically redundant—by new machines and novel forms of work organization. The value of artisanal skills was eroded by the rise of the factory system, and the value of skills in repetitive clerical and production tasks was eroded by computerization. Over the longer run, new forms of expertise gained value and catalyzed job creation, although often benefiting different workers from those who were displaced. Indeed, recent research estimates that more than 60 percent of current employment is found in occupational specialties that were not present in 1940. There is, however, substantial uncertainty about what types of new work will follow from the widespread use of AI, what skills it will require, what it will pay, how much of it there will be, and who will do it. But there is no question that AI will both strand some forms of human expertise and create demands for others.

The following questions are thus key to assessing the impact of AI advances on jobs:

- *What expertise will be substituted with or made obsolete by AI?* AI tools may soon equal or exceed human capabilities in a variety of tasks requiring elite expertise, such as digesting and summarizing large document collections; proofreading; writing certain business and legal documents; producing presentations and marketing materials, including charts, slides, and illustrations; and helping to manage complex systems such as computer networks and perhaps air traffic control systems. In such cases, AI is likely to substitute for human expertise, eroding the value of such expertise. Note that these types of expertise reflect knowledge work rather than physical work involving manipulating objects in the external world. At present, progress in robotics is slower than progress in AI that impacts the preceding types of knowledge work, but continued progress in robotics is anticipated in the longer term, over time substituting for routine and widely available physical expertise such as package delivery and food preparation in high-volume restaurants. The application of LLMs and other modern AI techniques may help speed up this progress.
- *What expertise will be augmented or newly demanded as a result of AI adoption?* The answer to this is even more uncertain than the answer to the previous

question. However, it appears likely that AI will be used in many (not all) cases to assist humans in performing tasks, thereby augmenting and complementing their expertise rather than substituting for it. For example, doctors are not expected to be replaced by medical diagnosis systems but rather to be assisted by them, resulting in an increase in the quality of decisions that will ultimately be made by the human doctor. Similarly, in many occupations AI tools will enable an increase in the quality of outcomes and the breadth of tasks human workers can undertake. For example, architects assisted by AI may be able to develop and evaluate a wider range of novel building designs. It is important to note that society has a choice in whether and where AI is used to augment human expertise versus substitute for it. Beyond augmenting existing types of human expertise, AI may lead to demand for new expertise as well. One broad possibility is that because AI will progress more rapidly in cognitive than physical applications, AI may raise the value of human dexterity and flexibility relative to cognitive capabilities. In this case, AI could prove more complementary to hands-on expertise, as seen in care work (including health care), the skilled trades, construction, and maintenance and repair. But this prediction presumes that AI capabilities in the cognitive realm substitute rather than complement labor in the cognitive domain and that the application of generative AI to robotics does not yield equally rapid progress.

- *How feasible will it be for workers to acquire newly valuable expertise?* Fortunately, this is a question over whose answer society has some control—one can choose whether to enrich the educational opportunities made available to the workforce and the degree to which governments subsidize the cost to workers of that retraining. Of course, the feasibility will depend on what types of expertise are newly required and the time and effort needed to acquire that expertise. Nevertheless, for many types of expertise, new training and certificate programs may provide a significant opportunity for governments to assist workers in adapting to the job disruptions likely to be produced from adoption of AI, while at the same time speeding its successful adoption and the attendant improvements to productivity and the GDP.
- *How will the organization of work and employer power dynamics evolve?* The consequences of previous technological transitions for the welfare of workers and the strength of the middle class have depended not only on the nature and application of technologies but also on the frameworks and legal institutions that have shaped their design, adoption, and use and the distribution of economic surplus among owners, managers, and line workers. The middle-class prosperity that accompanied the second industrial revolution is owed in



part to the success of labor unions and supporting legislation in negotiating for higher pay, reasonable hours, safe working conditions, and employment security. In more recent decades, U.S. worker bargaining power has eroded as labor union representation has shrunk, employers have mobilized against worker organizing efforts, and labor markets have “fissured” as companies outsourced tasks and responsibilities.<sup>5</sup> Many rights and norms are increasingly contestable in the era of AI, including the definition and ownership of intellectual property, the right to privacy, and expectations about surveillance and coercive monitoring. The outcome for worker and societal welfare necessarily depends on the legal, regulatory, and bargaining regimes in place. Absent well-functioning institutions, the committee does not presume that market outcomes will be socially desirable ones.<sup>6</sup> AI can also have other types of impacts on workers’ quality of life. For example, AI may be used to monitor worker activity and productivity: a potential tool to better link individual worker performance and compensation but also an enabler of surveillance that workers may find objectionable. For creative artists, AI raises a different set of issues. For example, if an AI system produces a song or movie in the voice, image, and style of a singer or actor, what compensation is that artist due? In short, AI adoption is likely to raise a variety of new questions about worker rights, privacy and surveillance, intellectual property, and more. Its impact is already far wider than simply shifting supply and demand for workers’ expertise.

## ARTIFICIAL INTELLIGENCE AND EDUCATION

Rising educational attainment in the workforce has played a key role in U.S. economic growth over the past 150 years, although large socioeconomic gaps in access to high-quality education over the past 50 years have contributed to rising income inequality. Access to high-quality primary and secondary education and to continuing education opportunities is likely to be a strong determiner of future U.S. economic growth.

Rapid advances in AI alter both the demand for and the supply of educational opportunities. As AI advances result in shifting skill demands for workers, access to

---

<sup>5</sup> D. Weil, 2014, *The Fissured Workplace: Why Work Became So Bad for So Many and What Can Be Done to Improve It*, Harvard University Press.

<sup>6</sup> This is a simple Coasean observation, not an indictment of markets per se: when the ownership of a property right (including the right to take or not take an action) is ill-defined, the market equilibrium set of actions involving the property right is likely to be inefficient because externalities in the exercise of this right will not be internalized. See R.H. Coase, 1960, “The Problem of Social Cost,” *Journal of Law and Economics* 3(October):1–44.

continuing education will be key to enabling the workforce to adapt. At the same time, AI may play a key role in providing new online learning environments for primary, secondary, and continuing education, especially environments that can customize to the differing learning needs and learning styles of each individual student. Indeed, multiple AI-enhanced online learning systems are already in widespread use.

Recent advances in LLMs offer the potential to design much more flexible, natural, and adaptive computer-based teaching environments than those currently in use, largely owing to their natural language and reasoning capabilities. Although there is not yet proof that these new LLM-driven methods lead to better student learning outcomes, they exemplify the explosion of creative new work going into designing and experimenting with this new generation of AI teaching tools.

Although the impacts of AI on the labor market remain uncertain, AI is likely to shift the demand significantly for different types of worker expertise and to result in a large increase in demand for continuing education and retraining programs to help workers acquire the expertise needed to adapt to the changing jobs environment. Furthermore, online tools that can help workers understand which job advances are within reach given their current skills and existing courseware can be instrumental in giving agency to workers as they try to adapt. The use of AI to build systems to fulfill these needs holds significant promise.

## MEASUREMENT

Given the great uncertainties in predicting exactly what technical advances in AI might occur in the near future and how these advances will impact demand for various types of expertise and workers, it is imperative to improve the observation and tracking of technical progress in AI, its adoption in practice, and its impacts on the workforce in near real time—and to share this information with the workforce.

Compared to 2017, when the prior National Academies' report was published, much more and better data are available to answer important questions about the impact of AI on the workforce. For example, the Census Bureau has created a new Annual Business Survey that includes questions about adoption of AI at the firm level, and this has been integrated with worker-level data housed by the Census Bureau. Furthermore, patent data from the U.S. Patent and Trademark Office have now been integrated into employer–employee data at the Census Bureau.

Private-sector data hold the potential to provide a much more real-time and large-scale picture of the state of job demand, skill supply, and salaries paid for different skill mixes, complementing the data collected by government agencies. For example,

LinkedIn holds a large real-time data set of résumés, ADP a large real-time data set of salaries paid, and Indeed.com a large real-time data set of job postings and salary offerings. Although companies such as LinkedIn and Lightcast make some statistical summaries of data available, challenges remain. Organizations are reluctant to share their detailed data owing to competitive concerns and privacy issues. Realizing the potential benefits of private-sector data will require new models for public–private data sharing, although the payoff in the ability to track and communicate changing skill demands to the workforce could be immense.

Many challenges remain in collecting and accessing the private data needed to ensure a clear picture of the state of AI adoption and of the corresponding changes in demand and supply of different worker skills. One challenge is to measure the complementary intangible capital investments and organizational restructuring that firms may undertake when adopting new technologies and that directly influence the workforce. Another is to modernize measures of productivity growth that were originally devised for goods-producing sectors such as manufacturing, to cover diverse productivity impacts of software tools. Yet another is to manage the different data schemas and formats used by different data sources.

A variety of steps can be taken to improve society’s ability to track and respond in a timely fashion to the impacts of advanced technologies on the workforce. New legislation could allow government statistical agencies to share their business data more widely. Steps can be taken to create shared data schemas across different government agencies, enabling greater integration of data from different government sources. Privacy-enhancing technologies such as multiparty secure computation can be adopted to enable greater data sharing while preserving privacy and confidentiality. New approaches that tap into the vast real-time data sets held by private organizations, while protecting commercial and privacy interests, could produce a major improvement in the ability of workers and decision makers to observe the current state and direction of demand for different skills and to react appropriately.

## FURTHER DATA COLLECTION AND RESEARCH INITIATIVES AND OTHER OPPORTUNITIES

As AI technology advances and is broadly adopted, there are numerous opportunities for industry, nonprofit institutions, worker organizations, academia, civil society, and government to influence the direction of this development.

## **Opportunities to Influence the Rates and Direction of the Development of Artificial Intelligence**

1. Support basic research in AI, including explainability, alignment, and safety research. Basic research is likely to have social returns that exceed the private returns, and thus markets will tend to underinvest in it.
2. Support research into standards and guardrails that could promote the adoption of AI in business environments.
3. Incentivize and establish enabling standards and regulations to encourage sharing and transparency regarding the data used to train advanced AI models, enabling a level playing field where new companies can enter the market and contribute to progress.
4. Fund AI research toward specific applications deemed to be of high priority to society, such as education and training, health care, climate change, and national security, where there are large positive externalities. Markets measure private benefits, not positive externalities experienced by parties who are not part of the transaction, even when such benefits are large. There is a strong economic rationale for government policies to promote research and activities in sectors with large positive externalities.
5. Support initiatives such as the National AI Research Resource and the Microelectronics Commons that can provide hubs for evolving research models and provide computational resources and foster talent needed to keep U.S. universities vital players in the development of frontier AI methods. Universities play critical roles both in educating the next generation of AI and other talent and in driving frontier research in AI and its applications across many fields.

## **Opportunities to Speed and Share the Productivity Benefits of Artificial Intelligence**

1. Invest in efforts by national statistical agencies and in public–private partnerships to collect and disseminate better data on AI adoption, AI’s contribution to productivity growth, and the factors that contribute to variation in this contribution across sectors and firms.
2. Support research into the effectiveness of policies that could enable labor mobility among occupations, firms, and geographical locations and help workers take better advantage of new job opportunities.
3. Support research into areas that contribute to regulatory uncertainty—such as product liability, copyright, privacy, and bias—and that complicate efforts of decision makers to assess benefits and risks, speed adoption and implementation, and drive productivity gains.

4. Support research to identify and assess the potential for AI technologies to create new harms either inadvertently or through abuse and help policy makers understand and consider the associated trade-offs and work with the private sector to develop sensible guardrails.
5. Support research to understand the implications for market concentration in AI, such as winner-take-most dynamics, and options for maintaining a competitive marketplace while still enabling the benefits of scale and scope.
6. Increase efforts to identify which specific occupational tasks are affected by AI as well as which old and new skills and expertise will therefore be in greater or less demand.
7. Support AI research that speeds scientific discovery, which is a key contributor to productivity growth.

### **Opportunities to Influence the Balance Among Worker Augmentation, New Work Creation, and Labor-Displacing Automation**

1. Support research on human-complementary AI and prioritize AI research that emphasizes improving quality through human–AI teaming.
2. Support research into best practices for fostering inclusive AI adoption within firms and organizations, including workers’ representative organizations and other ways to strengthen worker voice in business decisions.
3. Support research into how the cost of capital and the cost of labor affect business decisions about AI adoption for automation and for augmentation.
4. Support research to explore how workers might be protected by the establishment of new norms or legal protections for the rights of individuals to control and be compensated for the use of their likenesses, their other personal attributes, and their creative works.
5. Increase AI expertise within the federal government to support effective investment, oversight, and regulation across all mission areas, including transportation, labor, health care, education, environmental protection, public safety, and national security.
6. Evaluate and certify the quality of purported human-complementary technology before adopting it for publicly funded programs in such areas as education and health care.

## **Opportunities to Understand the Implications of Artificial Intelligence for Education and Assist Workers with Retraining and Continuing Education**

1. Support research on effective continuing education approaches, especially short-term programs that teach specific skills in high and growing demand—thereby helping foster workforce flexibility.
2. Support research on how AI, augmented reality, and other technologies can be used to improve education—in particular, the types of continuing education and retraining programs identified in the preceding item.
3. Support research into how standards and certification for training programs can help community colleges and other educational institutions more effectively signal graduates’ skills to employers and improve the match of new graduates to in-demand job opportunities.
4. Support an appropriate organization to develop, maintain, and disseminate a “career roadmap” that would enable workers to understand the shifting demand for different types of skills and workers as well as the continuing education opportunities available to them to acquire high-demand skills that will advance their careers.
5. Support research into new education objectives for all levels of education, including K–12, in order to provide the current and next generation with the knowledge and skills needed to take full advantage of future AI capabilities.

## **Opportunities for More Exact and Timely Tracking of Artificial Intelligence’s Impacts on the Workforce**

1. Improve and expand existing data collection efforts by government agencies, including high-frequency, real-time tracking of the use of AI by businesses and workers and the impact on the workforce.
2. Create new public–private data partnerships in which privately held data about skills supply and demand and wages currently paid for these positions are shared, and in which data about continuing education opportunities and their link to getting a better job are collected—with current summaries of both made available to the workforce, supporting workers’ efforts to improve their livelihoods.
3. Measure and mitigate disparate impact of new technologies on under-represented groups or communities as well as the global impact of differences in AI adoption between high-income and low-income countries.
4. Measure and characterize the heterogeneity in patterns of AI adoption across economic sectors, across firms within sectors, and across geographical regions, along with heterogeneous impacts on productivity and the workforce.

5. Undertake scenario planning to allow workers, businesses, and policy makers to gain a better understanding of the wide range of plausible scenarios for future advances in AI, including radical improvements.

\* \* \*

It is impossible to predict exactly the nature of the coming changes in AI and all of their effects on the economy and society. Accordingly, it makes sense to build in the ability for rapid data gathering and analysis to track these changes, and to build as flexible an approach as possible for reacting to the changes observed. In practice, this means increased capacity for research not only in AI but also in the social and behavioral sciences so that AI's implications can be better understood. It also means that rather than trying to predict any specific future path, society needs the flexibility to rapidly sense and respond to opportunities and challenges—and to be prepared for a variety of scenarios and possibilities. Most importantly, as AI becomes more capable, policy makers, business leaders, AI researchers, employers, and workers all have an opportunity to shape the future of the workplace and workforce in ways that are consistent with societal values and goals.