Congressional Budget Office
Nonpartisan Analysis for the U.S. Congress

Artificial Intelligence and Its Potential Effects on the Economy and the Federal Budget

DECEMBER | 2024

Summary

Artificial intelligence (AI) refers to computer systems that can perform tasks that have traditionally required human intelligence, such as learning and performing other activities that require cognitive ability. A general attribute of AI is its ability to identify patterns and relationships and to respond to queries that arise in complex scenarios for which the precise computational algorithm that is needed cannot be specified in advance.

Because AI has the potential to change how businesses and the federal government provide goods and services, it could affect economic growth, employment and wages, and the distribution of income in the economy. Such changes could in turn affect the federal budget. The direction of those effects—whether they increased or decreased federal revenues or spending—along with their size and timing, are uncertain. Some budgetary effects could occur relatively quickly, whereas others might take longer. In this report, the Congressional Budget Office provides an overview of the channels through which the adoption of AI could affect the U.S. economy and the federal budget.

How Might Artificial Intelligence Affect the Economy?

By increasing efficiency, enabling the development of new products, and altering the demand for workers, AI has the potential to change the economy, perhaps in ways that are difficult to predict. Whether or when those changes might occur is very uncertain. Surveys show that only 5 percent of businesses in the United States currently rely on AI to produce goods and services. For many businesses, customizing AI to their specific needs is costly, and it is unclear when those costs might fall. As a result, the use of AI is concentrated among larger, and younger, businesses in a few sectors of the economy—although that could change over time as profitable use of the technology became less dependent on a business's size.

Research into the performance of businesses that have implemented AI is still in its early stages, so conclusions from that research, which vary widely among studies, should be considered preliminary. So far, the research has found that businesses that implement AI can be expected to become more productive than businesses that do not. Extrapolation of those results to the broader economy suggests that if AI's use became more widespread, it would boost economic growth. Evidence for AI's impact on employment and wages is also sparse and varies by type of AI. Studies of generative AI indicate that it could enhance the productivity of low-skilled workers within a given occupation; studies of earlier forms of the technology have found that AI boosted the wages of some skilled workers.

How Might Artificial Intelligence Affect the Federal Budget?

The use of AI could affect the federal budget through two basic channels: the economy and the government. Within each channel, AI could have an impact on revenues and spending. The timing of budgetary effects may vary.

AI's Use in the Economy. The use of AI could affect the overall amount of income in the economy and its distribution among businesses, investors, and workers. An increase in income would, by itself, eventually push up federal revenues. Initially, however, revenues could decline as businesses deducted from their income the cost of initial investments in the technology. Moreover, because different categories of income are taxed at different rates, changes to how income is distributed among workers and businesses could alter federal revenues. In particular, depending on how the demand for workers shifted in response to the use of AI, tax receipts tied to labor income could rise or fall. For workers who were left permanently unemployed or who took lower-paying jobs because of businesses' adoption of new technology, income and payroll taxes could decline; however, workers who were made

Note: Unless this report indicates otherwise, all years referred to are federal fiscal years, which run from October 1 to September 30 and are designated by the calendar year in which they end.

more productive by AI could earn higher wages and remit larger tax payments. To the extent that AI created new kinds of tasks and jobs or led to economic growth through innovation, it could offset some potential losses of wages and taxes by increasing the demand for labor.

AI's use in the economy could change both mandatory spending (which does not require annual funding from the Congress) and spending subject to appropriation (which requires annual funding from the Congress). For instance, mandatory spending could increase to the extent that workers whose jobs were displaced by AI claimed benefits from federal income-support programs. But if AI boosted economic output and earnings, then such spending could decrease.

AI could also have an impact on federal spending through its use in the development of certain products, such as pharmaceuticals. For example, mandatory spending could increase to the extent that new drugs were paid for by federally subsidized health care programs. But that spending could decrease to the extent that the new drugs reduced the demand for other, more expensive, health care services. For AI's use in other types of health care that receive a federal subsidy and are provided by the private sector, the ultimate impact is often unclear. In general, if AI's use enabled individuals to live longer, healthier lives, then it could boost federal revenues and spending in the long term. Revenues would be higher if more taxes were paid over a longer period. Spending would be greater if more claims were made on Social Security, Medicare, and other programs by individuals who lived longer in retirement than they would have otherwise.

Spending subject to appropriation could increase if the Congress expanded funding for AI's continued development by the public and private sectors. That could happen through federal programs for research and development (R&D) as well as through programs that regulate the technology's use.

AI's Use by the Government. The government's use of AI could change the amount of revenues collected through taxes and other sources. Federal revenues could rise, for example, if the Internal Revenue Service (IRS) was able to use AI to bolster its auditing capability and taxpayers' compliance with the federal tax code. By contrast, revenues would decline if businesses or individuals were able to use AI tools to reduce the taxes they owed.

The government's use of AI could also have various effects on mandatory spending and spending subject to appropriation. In particular, successful use of AI to reduce fraud could result in fewer improper payments in the largest mandatory spending programs: Medicare, Medicaid, and Social Security. Those efforts could be undermined, however, if individuals used the technology to perpetrate fraud.

AI could affect the spending subject to appropriation of federal agencies that made use of the technology. Although investments in AI might initially increase spending, those costs could eventually decline, depending on the efficiencies that were realized. For example, if AI substituted for labor, staffing requirements could fall.

Development and Use of Artificial Intelligence

The past few decades have seen rapid strides in AI's capabilities. Today, AI is embedded in commonly used software (such as web browsers, phones, and home assistants) and found in more specialized applications (such as industrial robots and self-driving cars). Although recent advances in AI have fostered optimism that significant further progress is imminent, several factors could hinder the technology's near-term development and use. In particular, a lack of large datasets to train AI models could restrict improvements in accuracy and applicability. In addition, obstacles to obtaining the energy needed to power advanced AI systems could impede their use.

Advances in AI Systems

Machine learning techniques underlie most forms of AI today. Discriminative AI models, for example, can distinguish between different types of images, such as cancerous and noncancerous human cells, on the basis of the characteristics that they have been trained to associate with each type. Human intervention is necessary both to select the initial computational algorithm (or mathematical model) and training data and, ultimately, to assess the AI system's performance. The output of an AI model, then, is fundamentally a prediction of the correct response to a particular question or to a new set of conditions.

The most advanced AI systems are composed of neural networks that are made up of thousands or even millions of nodes where computation takes place. Deep-learning AI consists of dozens of layers of nodes, and information can be exchanged between individual nodes and between layers. A general attribute of such systems, which are typically trained on very large datasets, is their ability to identify relationships and respond to queries that arise in scenarios in which the needed computational algorithm cannot be specified in advance.

Generative AI systems can, upon receiving a question or prompt from users, provide answers and other types of content that are comparable to what would have been produced by a person. For example, large language models generate human-like text and perform natural language processing in which machines are able to understand and interact with human speech. Large language models rely on data collected from public internet sites, and they generate responses to user queries in the form of text or, where applicable, computer code. Those models are based on what is known as the transformer architecture, which is a deep-learning neural network that transforms an input sequence into an output sequence.[1] Other generative AI systems can produce user-specified images, video, and music—even in the style of specific artists, if desired.

The current capabilities of AI systems reflect several decades of dramatic advances. An early, and highly touted, application of AI was IBM's Deep Blue, a computer capable of playing chess. That system beat world chess champion Gary Kasparov in a series of matches in the late 1990s. That application of AI was relatively rudimentary in that it relied on substantial human input to learn the rules and various strategies of chess and to improve its performance. In 2016, a more advanced system, based on neural networks, defeated one of the world's top players of Go, which is a board game considered much more challenging than chess.

Other types of AI also have displayed rapid progress. For example, discriminative AI—which learns the boundaries between different classes of input data to identify (or categorize) them correctly—achieved a breakthrough at the ImageNet Large Scale Visual Recognition Challenge of 2012. In particular, by applying neural networks trained on a large dataset of images, one of the AI models competing in the challenge (AlexNet) achieved a very high degree of accuracy in recognizing images.[2] AI systems have continued to advance since then, and the performance of leading models on multitask tests has increased markedly over just the past few years.[3]

1. Ashish Vaswani and others, "Attention Is All You Need" (paper presented at the 31st Conference on Neural Information Processing Systems, December 2017, updated August 2023), https://arxiv.org/abs/1706.03762.

2. Olga Russakovsky and others, "ImageNet Large Scale Visual Recognition Challenge," *International Journal of Computer Vision*, vol. 115 (April 2015), pp. 211–252, https://doi.org/10.1007/s11263-015-0816-y.

3. "Multi-task Language Understanding on MMLU," https://tinyurl.com/2dk7jymj.

Near-Term Obstacles to Further AI Development

Despite recent rapid gains in the capability and accessibility of AI systems, the technology's history suggests that further progress could be interrupted. The field of AI research dates from the mid-1950s, and on more than one occasion, improvements in the technology heralded as breakthroughs have not led to follow-on success.

Recent advances in AI are attributed to the availability of large datasets and abundant power—in terms of both computer processing capability and the supply of energy—along with the development of increasingly complex computing algorithms. Some research suggests that AI could become widely used. However, in the near term, the adoption of AI also faces various obstacles.

One particular concern about wider adoption of AI is the looming scarcity of additional large datasets—containing either general, or task-specific, information—for training AI models. An additional concern is the ability of the electric power sector in the United States to supply sufficient electricity for powering and cooling the computers running the advanced AI software. The impact of data and energy scarcity is not certain; some analysts argue that although those factors exist, they need not constrain the technology's continued advancement.[4] This report focuses on the channels through which AI could affect the economy and the federal budget and does not assess the technology's future trajectory or the tasks that it might ultimately be able to perform.

Potential Effects of Artificial Intelligence on the Economy

Some people believe that AI could become as pervasive as electricity and computing are today. Although large investments in AI are currently being made by major businesses in the technology sector, its use by businesses overall remains limited. For that reason, research on the technology's economic effects is at an early stage, and the results from that research are uncertain. AI's effect on the performance of individual businesses is not yet well understood. Even though economy-wide productivity gains from the technology are expected, the size of those gains varies across studies. The general effects of AI's use on labor markets are also uncertain.

Productivity and Economic Growth

Many economists today view artificial intelligence as an emergent general-purpose technology. Such technologies

4. Jaime Sevilla and others, *Can AI Scaling Continue Through 2030?* (Epoch AI, August 20, 2024), https://tinyurl.com/bdnx4627.

do not have a sole definition, but they usually satisfy the following criteria:

- They can be applied throughout the economy;
- They are improved on a regular and sustained basis;
- Their use is accompanied by innovations in related areas (for example, new products and services); and
- They boost productivity and economic growth.

AI could transform society in the same way that technological advances like the steam engine and electrification did in the distant past and as computing and the internet have done over the past few decades.[5] For example, similar to the wide applicability of information technologies, AI has been found to boost the productivity of researchers looking to create new products in a variety of disciplines, such as materials science.[6]

Businesses' current use of AI remains limited, though. Surveys show that only 5 percent of businesses of a broad range of sizes in the United States (accounting for 9 percent of employment) currently incorporate AI—in more than an incidental way—in their production of goods and services. Furthermore, those businesses that use AI tend to be found in particular industries: Businesses in the "information" and "professional, scientific, and technical services" industries are roughly twice as likely as other businesses to be using AI.[7] Businesses' adoption of AI is found to be more frequent when surveys do not require respondents to identify the technology's use as being a significant part of the production process. For example, one survey found that about 28 percent of individuals reported using generative AI simply "for (their) job."[8]

Research into the economic impact of AI, in terms of the productivity gains of businesses implementing it as well as the implications of those gains for the economy's growth, is still relatively new. A common approach taken by researchers to study those effects has two steps. First, researchers identify tasks that AI is likely either to take over from workers altogether or to enable them to do better. Then, researchers convert those efficiency gains (or reductions in cost) to increases in the amount of a business's output that can be produced per hour worked by the business's employees—known as labor productivity. (Researchers sometimes use an alternative performance measure known as total factor productivity, which is the quantity of output produced relative to the amount of all inputs into production.)

Although the research on AI's economic effects focuses on labor, AI can make physical capital more productive as well. One example would be by enabling robots to grasp and manipulate workpieces that they have not encountered before. That type of change would improve efficiency because processes would no longer need to be programmed each time a workpiece changed shape or some other dimension. As a result, the cost of manufacturing disparate products could fall.

According to research, businesses that implement AI can typically be expected to be more productive than those that do not. Some studies extrapolate from those impacts to estimate how much the economy's output could rise if AI's use became more widespread. (Such findings cannot directly be interpreted as having a similar effect on productivity in the economy as a whole. To be done properly, studies must take into account the interaction effects between, say, a new technology and other variables of interest, and their broader economic effects; such interactions are difficult to predict.) Although the impact of AI on the performance of individual businesses or of

5. Iain M. Cockburn, Rebecca Henderson, and Scott Stern, *The Impact of Artificial Intelligence on Innovation*, Working Paper 24449 (National Bureau of Economic Research, March 2018), www.nber.org/papers/w24449; and Timothy F. Bresnahan and M. Trajtenberg, "General Purpose Technologies: 'Engines of Growth'?" *Journal of Econometrics*, vol. 65, no. 1 (January 1995), pp. 83–108, https://doi.org/10.1016/0304-4076(94)01598-T.

6. Aidan Toner-Rodgers, *Artificial Intelligence, Scientific Discovery, and Product Innovation* (Massachusetts Institute of Technology, November 6, 2024), https://aidantr.github.io/files/AI_innovation.pdf.

7. Kathryn Bonney and others, *Tracking Firm Use of AI in Real Time: A Snapshot From the Business Trends and Outlook Survey*, Working Paper 32319 (National Bureau of Economic Research, April 2024), pp. 2–4 and p. 39, www.nber.org/papers/w32319. In that survey, businesses were asked whether they used AI "in producing goods or services," and the following examples were provided as guidance for respondents: "machine learning, natural language processing, virtual agents, voice recognition, etc." Hence, reported AI adoption should not reflect incidental use, such as the AI embodied in common software like web browsers or business applications.

8. Alexander Bick, Adam Blandin, and David J. Deming, *The Rapid Adoption of Generative AI*, Working Paper 32966 (National Bureau of Economic Research, September 2024), p. 11, www.nber.org/papers/w32966.

the economy overall is expected to be positive, the size of that impact varies greatly among studies.[9]

The conclusions of current studies about AI's broad economic impact are preliminary. In those studies, actual outcomes are not observed. Instead, the studies attempt to link AI to specific tasks, tasks to workers, and workers affected by AI to changes in businesses' performance. Surveys of businesses that are *not* currently using AI suggest that the eventual use of AI and its impact on workforces and performance remain poorly understood. For example, of the businesses that report that they do not use AI, four out of five state that "AI is not applicable to this business."[10]

Expense is a major obstacle to greater use of AI. Although leading AI systems have made great strides in performance in recent years, the cost of training them has increased—from tens of millions of dollars to hundreds of millions of dollars. Some analysts project AI training costs to reach $1 trillion by the end of the decade.[11] Customizing a given model so that it can be applied by individual businesses can entail costs that make AI prohibitively expensive for many businesses, even if that customization involves relatively established and advanced types of AI, such as computer vision (which allows a computer to analyze and identify images).[12]

Neither the lack of evidence to date that AI is substantially changing how goods and services are produced nor the fact that most businesses report that AI is not now advantageous to them means that the technology will not eventually have significant economic effects. For example, some researchers argue that AI will ultimately reach the stage of artificial general intelligence, at which point it will be able to carry out cognitive-based tasks as well as, or better than, human beings. If that occurred, the economic implications could be significant—in particular, for labor markets and the distribution of income.[13]

Indeed, a substantial lag often follows the availability of important new technologies and their measurable impact on the economy. In the case of electrification, for example, four decades after the first central power station opened for business in the United States, only about half of the mechanical drive capacity in factories had been electrified.[14] That said, in terms of technological diffusion, it is difficult to pinpoint where AI might be relative to such earlier timelines.

Employment and Wages

Research on the impact that AI has had on employment and wages is sparse. One survey found that roughly nine out of 10 businesses using AI report that they have not changed the number of workers they employ as a result and have no plans to do so. Among businesses that have made such changes or are planning to, employment increases are as likely to occur as decreases.[15]

The most accurate way to gauge AI's likely impact on the workforce is by assessing how the technology will affect the tasks that workers perform. One recent study found that the eventual impact of AI—specifically, large language models—on workers' jobs would vary: 80 percent of the U.S. workforce could have at least one-tenth of their tasks

9. Daron Acemoglu, *The Simple Macroeconomics of AI*, Working Paper 32487 (National Bureau of Economic Research, May 2024), www.nber.org/papers/w32487; Michael Chui and others, *The Economic Potential of Generative AI: The Next Productivity Frontier* (McKinsey, June 2023), https://tinyurl.com/y9tevm2n; and Joseph Briggs and Devesh Kodnani, *The Potentially Large Effects of Artificial Intelligence on Economic Growth* (Goldman Sachs Economics Research, March 26, 2023), https://tinyurl.com/4anjjbdf.

10. Kathryn Bonney and others, *Tracking Firm Use of AI in Real Time: A Snapshot From the Business Trends and Outlook Survey*, Working Paper 32319 (National Bureau of Economic Research, April 2024), Table 7, p. 37, www.nber.org/papers/w32319.

11. Anton Korinek, *Economic Policy Challenges for the Age of AI*, Working Paper 32980 (National Bureau of Economic Research, September 2024), p. 5, www.nber.org/papers/w32980; and Ben Cottier and others, *The Rising Costs of Training Frontier AI Models* (arXiv, May 31, 2024), https://arxiv.org/abs/2405.21015.

12. Martin Fleming, Wensu Li, and Neil C. Thompson, *The Last Mile Problem in AI: Why Job Automation Will be Slower Than Technological Progress Suggests* (Brookings, August 29, 2024), www.brookings.edu/articles/the-last-mile-problem-in-ai/.

13. See, for instance, Anton Korinek, *Economic Policy Challenges for the Age of AI*, Working Paper 32980 (National Bureau of Economic Research, September 2024), www.nber.org/papers/w32980.

14. Paul A. David, "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox," *American Economic Review*, vol. 80, no. 2 (May 1990), pp. 355–361, www.jstor.org/stable/2006600.

15. Kathryn Bonney and others, *Tracking Firm Use of AI in Real Time: A Snapshot From the Business Trends and Outlook Survey*, Working Paper 32319 (National Bureau of Economic Research, April 2024), Table 6, p. 36, www.nber.org/papers/w32319.

affected by AI, and 19 percent of workers could see at least half of their tasks affected.[16]

That potential scope of the use of AI makes it difficult to estimate how overall employment might change as a result of the technology and, in particular, how many workers might be displaced by it. Quantifying those effects would require establishing a threshold level of tasks for which AI is preferred over a worker and beyond which that worker would be deemed superfluous and have their job terminated. That threshold might not be the same for all businesses, because the increase in demand spurred by the decline in products' prices brought on by AI-driven cost reductions could vary. In some cases, greater demand could lead to retaining—and potentially adding—workers who would perform fewer tasks at a higher volume of production; in other cases, any additional demand could be insufficient to warrant keeping on staff the employees most affected by AI.[17]

The technology's effect on employees' tasks can take the form of either substituting for workers in accomplishing a particular task or serving as a complement to them—that is, make them more productive. One task-based study argues that the substitution effect of AI will outweigh the complementarity effect; as a consequence, AI would reduce labor's share of production.[18]

Other research examines the substitution and complementarity effects of AI on workers of different skill levels. Those studies explore how the wage differential between low-skilled and high-skilled workers could change based on a business's use of AI. If AI substituted for high-skilled workers, then the wage differential between the two groups of workers would decline; if it complemented high-skilled workers, the wage differential would rise. The opposite would be true for AI's impact on low-skilled workers; in particular, if the technology complemented those workers, the wage differential would fall.[19]

Evidence shows that generative AI can serve, at least to some degree, as a complement to low-skilled workers within a given occupation. By contrast, research on earlier forms of AI has found that the technology boosted the wages of some skilled workers.[20] A study of generative AI's impact on the productivity of customer support agents—measured by the number of issues resolved per hour—found that AI increased the productivity of entry-level and low-skilled agents by 34 percent. In contrast, experienced and highly skilled workers did not show significant gains in productivity. The authors suggest that "the AI model disseminates the best practices of more able workers and helps newer workers."[21] Other analysts argue that AI may be similarly skill-enhancing elsewhere, enabling workers who lack the experience or expertise of higher-paid employees to take on greater responsibility in the workplace and thus dampening earnings inequality among workers.[22]

Potential Effects of Artificial Intelligence on the Federal Budget

AI could affect the federal budget through two channels: its use in the economy, and its use by the government. Within each channel, AI could have an impact on revenues, mandatory spending, and spending subject to appropriation. Some budgetary effects might occur relatively quickly, whereas others might take longer to show up.

Through its use in the economy, AI could affect revenues by changing the amount of national income and its distribution. The technology could affect spending

16. Tyna Eloundou and others, *GPTs Are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models*, Working Paper 10130 (arXiv, March 2023, updated August 2023), https://arxiv.org/abs/2303.10130.

17. Other approaches to quantifying the effects that AI could have on employment consider how it could take the place of the "expertise" that workers supply to their employer and posit how the demand for different types of labor will vary depending on the extent to which AI substitutes for the expertise associated with it. See the National Academies of Sciences, Engineering, and Medicine, *Artificial Intelligence and the Future of Work* (2024), pp. 88–95, https://tinyurl.com/4cxu6uwy.

18. Daron Acemoglu, *The Simple Macroeconomics of AI*, Working Paper 32487 (National Bureau of Economic Research, May 2024), www.nber.org/papers/w32487.

19. David E. Bloom and others, *Artificial Intelligence and the Skill Premium*, Working Paper 32430 (National Bureau of Economic Research, May 2024), www.nber.org/papers/w32430; and Mauro Cazzinga and others, *Gen-AI: Artificial Intelligence and the Future of Work*, IMF Staff Discussion Note SDN2024/001 (International Monetary Fund, January 2024), https://doi.org/10.5089/9798400262548.006.

20. Edward W. Felten, Manav Raj, and Robert Seamans, *The Occupational Impact of Artificial Intelligence: Labor, Skills, and Polarization* (NYU Stern School of Business, September 8, 2019), https://ssrn.com/abstract=3368605.

21. Erik Brynjolfsson, Danielle Li, and Lindsey R. Raymond, *Generative AI at Work*, Working Paper 31161 (National Bureau of Economic Research, revised November 2023), www.nber.org/papers/w31161.

22. David Autor, *Applying AI to Rebuild Middle Class Jobs*, Working Paper 32140 (National Bureau of Economic Research, February 2024), www.nber.org/papers/w32140.

by altering participation in means-tested programs and the use of federally subsidized health care services, as well as by inducing more funding for federal programs that could support AI's continued development and governance.

Through its use by the federal government, AI could affect both revenues and spending by increasing the efficiency of the government in collecting tax revenues and in distributing those revenues through transfer payments. AI also could enable improvements in the goods and services provided by the government, spurring federal programs to spend more to take advantage of the technology. Overall, the ultimate impact of AI on federal revenues and spending is uncertain.

Budgetary Impact of AI's Use in the Economy

The use of AI in the economy could affect revenues by changing the amount and distribution of income. In addition, it could affect mandatory spending, by changing participation in means-tested programs and the use of federally subsidized health care services, and spending subject to appropriation, by changing the policy choices that underlie funding decisions for various programs each year.

Revenues. The amount of income and how that income is distributed could change as a result of AI's use in the economy. An increase in income that affected businesses, investors, and workers proportionately would boost federal revenues. Because different categories of income are taxed at different rates, though, changes to the distribution of income among categories could offset that increase. Moreover, businesses' taxable income could decline, at least temporarily, as a result of investments in AI.

If businesses became more productive, profits would eventually increase, pushing up earnings from capital, which accrue either to the owners of privately held businesses or to investors in publicly traded ones. As a result, tax payments on business income and returns from equities (in the form of dividend payments or capital gains) would rise. Those higher tax revenues could follow a period during which tax payments associated with AI fell because of tax deductions taken for investments in it and, potentially, because of lower profits during the initial (and typically costly) stage of implementing a new technology.

The effect of AI on taxable labor income is uncertain. It would depend on the extent to which the positive effects of a larger economy were offset by the potential negative effects that could occur if AI substituted for labor. Tax receipts tied to labor income could rise or fall depending on how the demand for workers shifted in response to the use of AI. If the use of AI was complementary to existing jobs (rather than a substitute for them), it could enable workers to do their current jobs better and perhaps undertake new tasks as well—making employees more productive and leading to higher wages. AI could also spur job creation if it enabled the production of new goods and services. Furthermore, if AI's use led to income gains from higher wages and job creation, it could have a positive impact on federal receipts tied to labor income.[23] In contrast, if workers were left permanently unemployed or were reallocated to lower-paying jobs by the technology, income and payroll taxes would decline.

Mandatory Spending. Participation in means-tested programs could rise or fall depending on the net effect that AI had on employment. Like revenues, mandatory spending programs—those whose spending is generally determined by formulas and eligibility criteria established by lawmakers rather than by annual appropriations—could also be affected by economic factors.

For example, even if the workers displaced by AI eventually found new jobs, that labor reallocation process could be lengthy. As a result, spending could increase over an extended period for federal income-support programs (including unemployment benefits, health care subsidies, and other means-tested programs), which provide cash payments or other assistance to people with relatively low income and few assets. But if AI increased economic output, employment, and wages, then mandatory spending on income-support programs could fall.

Changes in mandatory spending could also occur as a result of the use of specific types of AI in the economy. For example, AI is currently being applied in pharmaceutical R&D.[24] If patients' consumption of the new drugs that AI helped discover were paid for by health insurance subsidized by the federal government, then federal outlays to provide pharmaceutical benefits could increase.

23. Tania Babina and others, "Artificial Intelligence, Firm Growth, and Product Innovation," *Journal of Financial Economics*, vol. 151 (January 2024), https://doi.org/10.1016/j.jfineco.2023.103745; and Dean Alderucci and others, "Quantifying the Impact of AI on Productivity and Labor Demand: Evidence From U.S. Census Microdata" (draft, 2019), https://api.semanticscholar.org/CorpusID:237265713.

24. Government Accountability Office, *Artificial Intelligence in Health Care: Benefits and Challenges of Machine Learning in Drug Development*, GAO-20-215SP (January 21, 2020), www.gao.gov/products/gao-20-215sp.

Box 1.

Artificial Intelligence and Federally Subsidized Health Care

Artificial intelligence (AI) could be used in a number of areas of federally subsidized health care. Although the government would be providing financial support to patients, say, through the federal tax subsidy for employer-sponsored health insurance or through the Medicare and Medicaid programs, service providers would typically be in the private sector. For this analysis, the Congressional Budget Office would consider such cases to be examples of the use of AI in the economy rather than of the use of AI by the federal government—even though the government would be subsidizing a sizable share of the cost. In contrast, an example of health care provided by a federal agency is services delivered by the Veterans Health Administration, which is funded through annual appropriation acts.

In many cases, the net effect on the federal budget of the use of AI in the economy is unclear. That is in part because AI can be used by participants in the health care system who have different objectives with potentially countervailing effects. (The text of this report discusses similar outcomes for other applications of AI.) Furthermore, the effects of AI's use in health care may vary across applications, sometimes leading to increases in federal spending and other times, to reductions.

One potential application of AI in the health care sector is for prior authorization (PA), a cost-control process that insurers use to limit access to high-cost services and drugs. Insurers could use the technology to streamline reviews of PA submissions, making it cheaper for them to roll out broader PA plans. Although broader PA plans would reduce health care costs and federal subsidies, health care providers could, in turn, use AI to clear the hurdles put in place by insurers, which would have an offsetting effect on those savings.

Another potential application of AI is to streamline the provision of health care services. By introducing efficiencies in how health care practices are managed and administered, AI could enable providers to furnish more services. In that case, costs to the federal government would increase. However, by enabling greater machine- and software-based health monitoring and treatment, AI could lower the cost (and, potentially, the allowable reimbursement) of those services and thus reduce federal spending.

The provision of health care services is a notable example of the potential difficulty in distinguishing between AI's use in the economy and its use by the federal government. That is because federal funding for activities that use AI can extend beyond the government to the private sector. For example, the National Institutes of Health (NIH) could fund AI-supported R&D that is carried out by staff in its own laboratories and by researchers elsewhere. In the former scenario, government employees would use AI, a clear-cut instance of use of the technology by the government. In the latter scenario, the technology could be used by, say, university faculty and affiliated researchers whose jobs were not dependent on federal financial support—even though some of the projects reliant on AI were contingent on receiving federal funds. A similar overlap applies to AI's potential use in other types of health care that benefit from federal subsidies (see Box 1).

Spending Subject to Appropriation. AI could affect spending for a variety of federal programs that require an appropriation by the Congress each year. For example, the reallocation of employment brought about by AI's use in the economy might lead lawmakers to provide larger appropriations for educating and retraining displaced workers.

The use of AI in the economy could also influence funding for federal programs unrelated to employment. Businesses that produced AI systems might invest in R&D to advance the technology, for instance. If they were not able to appropriate all of the returns to that R&D, they would tend to invest less than the economically efficient level (that is, the level at which the financial payoff from additional R&D matched the investment in it). That underinvestment could be more likely to occur for AI because the technology is expected to foster innovation in many sectors of the economy, thus making the value of R&D for different users difficult to determine in advance. As a result, the Congress may supplement private-sector R&D efforts by increasing federal funding for research and development into AI.

Lawmakers have enacted legislation that addresses AI's use both in the economy and by the federal government. Examples include the following:

- The AI in Government Act of 2020 (Public Law 116-260), which created within the General Services Administration the AI Center of Excellence to facilitate adoption of AI by federal agencies;

- The National Artificial Intelligence Initiative Act of 2020 (P.L. 116-283), which supports AI's use in both the public and private sectors; and
- Other laws, such as the Advancing American AI Act (P.L. 117-263), which focuses on monitoring and setting policies for AI's use by the government.

As the technology is increasingly applied in the private sector, lawmakers might enact additional legislation that could affect federal spending. In particular, widespread use of AI might raise additional governance issues for the Congress and lead it to fund new or expanded federal agencies. One potential issue concerns data privacy. The most advanced AI systems rely on large quantities of data, which could result in improper use of personal information.

Budgetary Impact of AI's Use by the Federal Government

The federal government's use of AI could affect both revenues and spending. The overall effect of AI on revenues collected by the tax system is ambiguous: Some factors will probably increase revenues, and others will decrease them. In terms of mandatory spending, AI could reduce the amount of payments made by the federal government, thus decreasing spending. But individuals' use of AI could counteract the government's efforts—if, for example, people used AI to forge identity documents that allowed them to fraudulently claim benefits. By enabling improvements in the quality of the goods and services that the government provides and the efficiency with which the government provides them, AI could lead to changes in mandatory spending and in the spending subject to appropriation of federal programs that sought to take advantage of the technology.

Revenues. One way that AI could increase revenues is by improving the ability of the Internal Revenue Service to detect noncompliance by taxpayers and enforce laws intended to ensure compliance. The IRS is using AI to improve its estimates of the federal tax gap, which is the difference between the tax payments that individuals voluntarily make on a timely basis and the amount that they owe. Estimates of the tax gap are based on an analysis of audited tax returns that identifies noncompliance with the tax code missed by IRS auditors. The IRS expects that AI will enable examiners to identify noncompliance better.[25] If the agency was able to use that information to increase compliance with tax laws, revenues would increase.

AI could be used to reduce revenues as well. One way that might happen is if businesses and corporations used AI to reduce their tax liability. Some areas of tax law are particularly complex, and AI could be used to legally decrease tax payments by identifying aspects of the tax code that a person might miss.

Mandatory Spending. The government's use of AI could have various effects on mandatory spending. One notable example concerns improper payments, which are payments that should not have been made or that were made in the incorrect amount. Those payments, which cost the government billions of dollars each year, can arise from incorrect and fraudulent billing in federal programs that pay private entities to provide services. By analyzing billing data and determining the characteristics of reimbursement claims that are likely to be incorrect or fraudulent, AI could help identify and reduce unwarranted federal payments and, as a result, lower spending. Efforts are already underway to apply AI to reduce improper payments in the Medicare, Medicaid, and Social Security programs.[26]

Estimates of improper payments vary substantially in scope, time period, and amount. The Government Accountability Office (GAO) reports that executive branch agencies' cumulative estimates of the improper payments they made from 2003 to 2023 totaled about $2.7 trillion, or roughly $130 billion on an annual basis (not adjusted for inflation).[27] Adopting a different methodology, which applies a statistical model to 12 federal agencies' spending and is limited to estimating excessive payments solely from fraud, GAO found that

25. Government Accountability Office, *Tax Gap: IRS Should Take Steps to Ensure Continued Improvement in Estimates*, GAO-24-106449 (June 5, 2024), www.gao.gov/products/gao-24-106449.

26. See the entries for the Department of Health and Human Services—Centers for Medicare & Medicaid Services and the Social Security Administration in the database of uses of AI by federal agencies at AI.gov, "The Government Is Using AI to Better Serve the Public" (September 1, 2023), https://ai.gov/ai-use-cases/.

27. Government Accountability Office, *Improper Payments: Information on Agencies' Fiscal Year 2023 Estimates*, GAO-24-106927 (March 26, 2024), pp. 1–2, www.gao.gov/products/gao-24-106927.

federal losses from 2018 to 2022 totaled $233 billion to $521 billion—or about $47 billion to $104 billion annually over that five-year period—and that the range of losses was attributable both to uncertainty surrounding the estimates and to the different risk conditions prevailing over the period.[28]

For several reasons, it is unclear whether or how much the use of AI might eventually reduce the improper payments made by the federal government. Evidence about the technology's effectiveness in the applications described above is lacking, and it is unclear how widely AI will eventually be deployed in those applications. (Studies of states' use of the technology to establish eligibility for public assistance programs point to shortcomings in AI's performance.)[29] Furthermore, GAO suggests that although AI could be a useful tool for the federal government's efforts to reduce fraud, the technology could also be used by individuals to perpetrate fraud. For example, AI could be used to create fake images for falsified documents.[30]

As noted earlier, the use of AI by federal employees carrying out R&D in health care could cause mandatory spending to rise if the outcome of that research, such as demand for new pharmaceuticals, received a federal subsidy. That increase could be offset by a decline in spending elsewhere if the new drugs eliminated the need for other types of federally subsidized health care.

Spending Subject to Appropriation. For federal programs that make use of the technology, AI could affect their spending subject to appropriation. AI could, for example, be especially useful at the Department of Defense, an agency that relies heavily on information analysis.[31] DoD has a large workforce—including members of the military, federal civilian workers, and private contractors—and many complex tasks that it must undertake to carry out its mission. DoD could use AI to manage its operations, enhancing the capability of current weapon systems and developing new ones. Effects on costs could be mixed: Increases in operational efficiency would tend to reduce costs; developing new weapon systems could either increase or decrease costs depending on how the cost of new systems compared with the cost of existing ones. And the AI systems themselves would require spending by the government. Less than 1 percent of DoD's 2024 budget request is for AI.[32]

For other programs and agencies, AI could also have an impact on spending subject to appropriation. In particular, the increased labor productivity projected from AI's use in the private sector could be realized for the federal workforce or federal contractors through efficiencies (and cost reductions) in how the federal government operates or delivers services. One example is AI's support for, and potential replacement of, staff in the call centers of programs operated by the General Services Administration, Internal Revenue Service, and Department of State.

Although investments in AI might initially increase agencies' spending, costs could eventually decline. The size of any decrease would depend on how many efficiencies were realized. For example, if AI substituted for labor, it could reduce staffing requirements.

Other Considerations About Budgetary Effects

Many aspects of AI's adoption, use, and effects are uncertain. In assessing the ways that AI could affect federal revenues and spending, several considerations should be kept in mind.

First, the examples in this report are illustrative. Whether describing projects that are underway or suggesting projects that are possible, the examples this report provides are intended simply to illustrate the various ways that the technology could affect federal revenues or spending. The

28. Government Accountability Office, *Fraud Risk Management: 2018–2022 Data Show Federal Government Loses an Estimated $233 Billion to $521 Billion Annually to Fraud, Based on Various Risk Environments*, GAO-24-105833 (April 16, 2024), www.gao.gov/products/gao-24-105833.

29. Kevin De Liban, *Inescapable AI: The Ways AI Decides How Low-Income People Work, Live, Learn, and Survive* (TechTonic Justice, November 2024), www.techtonicjustice.org/reports/inescapable-ai.

30. Government Accountability Office, *Fraud Risk Management: 2018–2022 Data Show Federal Government Loses an Estimated $233 Billion to $521 Billion Annually to Fraud, Based on Various Risk Environments*, GAO-24-105833 (April 16, 2024), p. 35, www.gao.gov/products/gao-24-105833.

31. Department of Defense, *Summary of the 2018 Department of Defense Artificial Intelligence Strategy: Harnessing AI to Advance Our Security and Prosperity* (February 12, 2019), https://tinyurl.com/mrxpv6tm.

32. Office of the Under Secretary of Defense, Comptroller/Chief Financial Officer, *United States Department of Defense Fiscal Year 2024 Budget Request* (March 2023), p. 16, https://tinyurl.com/mr26n622.

federal government may currently, and in the future, use AI in many more ways than the applications listed above, but the extent of the technology's eventual adoption and use are uncertain.[33]

Second, outlays for AI could take the place of other federal spending. For statistical analysis, for example, AI could replace software that the government has been using and would continue to rely on if AI was not available. Use of AI might also entail greater reliance on cloud computing, which could take the place of upgrades to agencies' computer hardware. In that case, determining the true budgetary cost of AI would require netting out the previous federal spending that was done to pay for licenses and other costs associated with using the abandoned software and hardware.

Third, the impact of AI's use could manifest in different ways over time, and those ways could have countervailing effects on the federal budget. If AI improved the quality of federally subsidized pharmaceuticals and health care more broadly, it could enable individuals to live longer, healthier lives. That development could increase revenues (because more taxes would be paid over a longer period) as well as spending (because benefits would be provided through Social Security, Medicare, and other programs for more years) in relation to the amounts that would otherwise have been received and spent.[34]

The budgetary effects would also depend on whether federal agencies made effective decisions about adopting AI systems. For example, AI models could be expensive to acquire, train, and operate. At a minimum, those costs would cover buying the software, paying employees to maintain and operate the AI models, and purchasing the electricity needed to run them; costs would probably vary by application and by agency. For AI systems to either increase revenues or reduce spending, those positive budgetary impacts would need to outweigh the negative effects of the outlays required to obtain and implement the technology.

Fourth, AI could have impacts that would not be fully reflected in federal budget totals. For government spending on pharmaceutical R&D, for instance, AI could enable innovations of higher quality than would have been possible otherwise. Consequently, even if new drugs developed with AI were more expensive, they may deliver more benefits to patients per dollar of R&D spending than would have been possible without the technology. In such cases—which could also arise from other, nonpharmaceutical applications of AI, such as increased security through a more robust national defense—the payoff from greater cost-effectiveness would not necessarily be captured, or at least identifiable, in data about federal revenues or spending.

33. See AI.gov, "The Government Is Using AI to Better Serve the Public" (September 1, 2023), https://ai.gov/ai-use-cases/. That database lists specific instances of the use of AI by the federal government. Also, agencies implementing AI typically have a portion of their website dedicated to it.

34. Government Accountability Office, *Artificial Intelligence in Health Care: Benefits and Challenges of Technologies to Augment Patient Care*, GAO-21-7SP (November 30, 2020), www.gao.gov/products/gao-21-7sp.

This report was prepared at the request of the Chairman of the House Committee on the Budget. In keeping with CBO's mandate to provide objective, impartial analysis, this report makes no recommendations.

Nathan Musick wrote the report with contributions from Tess Snyder (a dissertation fellow at CBO) and with guidance from Joseph Kile. Leigh Angres, Aaron Betz, Sheila Campbell, Nicholas Chase, Devrim Demirel, Ann E. Futrell, Sebastien Gay, Tamara Hayford, John McClelland, Jaeger Nelson, Kevin Perese, Asha Saavoss, Molly Saunders-Scott, Joshua Shakin, Chad Shirley, and Chapin White provided comments.

Kevin Bryan of the University of Toronto commented on an earlier draft. The assistance of an external reviewer implies no responsibility for the final product; that responsibility rests solely with CBO.

Mark Doms, Mark Hadley, and Jeffrey Kling reviewed the report. Christine Bogusz edited it, and R. L. Rebach prepared the text for publication. This report is available at www.cbo.gov/publication/60774.

CBO seeks feedback to make its work as useful as possible. Please send comments to communications@cbo.gov.

Phillip L. Swagel
Director