October 2024

Working Paper

The fiscal frontier: Projecting AI's long-term impact on the US fiscal outlook

Benjamin H. Harris, Neil Mehrotra, and Eric So

This working paper is available online at: https://www.brookings.edu/center/center-on-regulation-and-markets/

Center on Regulation and Markets at BROOKINGS

The Center on Regulation and Markets at Brookings creates and promotes rigorous economic scholarship to inform regulatory policymaking, the regulatory process, and the efficient and equitable functioning of economic markets. The Center provides independent, non-partisan research on regulatory policy, applied broadly across microeconomic fields.

Disclosure

The views here represent those of the authors and do not reflect the views of the Federal Reserve Bank of Minneapolis or the Federal Reserve system.

The Brookings Institution is financed through the support of a diverse array of foundations, corporations, governments, individuals, as well as an endowment. A list of donors can be found in our annual reports published online here. The findings, interpretations, and conclusions in this report are solely those of its author(s) and are not influenced by any donation.

The Fiscal Frontier: Projecting AI's Long-Term Impact on the US Fiscal Outlook

Benjamin H. Harris, Neil Mehrotra, and Eric So

October 2024

Abstract:

We simulate the impact of artificial intelligence (AI) on the long-term federal fiscal outlook. This paper introduces a framework for how AI will affect fiscal budgets through four primary channels: mortality rates and the size of the population, the price of health care services, demands for health care services, and aggregate productivity. Using this framework, we show that the nature of the AI shock is critical, as the impact of the shock on annual budget deficits could range from an increase of 0.9 percent of GDP to a decrease of 3.8 percent of GDP, with the latter instance effectively halving annual budget deficits.[1]

[1] We thank Liam Marshall for outstanding research assistance. We also thank Emilia Javorsky for a helpful discussion at the Brookings Artificial Intelligence Author's Conference, and other participants in the Brookings Artificial Intelligence Author's Conference for helpful comments and feedback. All errors or omissions are our own.

1. Introduction

Human history is marked by repeated waves of technological innovation and progress, with each wave reshaping civilization in its wake. From the mastery of electricity and the invention of the wheel to the printing press and the steam engine, each era-defining technological advancement has expanded the horizons of human capability. Not only does each of these waves add to society's toolkit, but they also profoundly influence culture, society, and individual lives. Now, in the 21st century, we sit at the crest of another wave, marked by advancements in artificial intelligence (AI).

In recent years, AI has transformed from a niche tool to a mainstream application. This transformation has been hailed by techno-optimists who envision a future in which AI not only enhances productivity and economic efficiency but also solves complex societal problems. With each algorithmic breakthrough and successful application, the promise of AI reinforces the hope that we might overcome some of the most daunting challenges facing humanity. This sentiment is widely shared by many who are witnessing the formidable strides made by AI. Our study is motivated by the fact that, amidst the latest wave of techno-optimism fueled by AI, virtually no attention is being paid to how AI might shape the fiscal outlook.

The impact of AI on federal spending and revenues is highly uncertain, given the nascent evolution of AI and the unpredictable impact of AI on economic activity. From the outset, we acknowledge this uncertainty and aim to model a series of representative shocks that provide a fulsome representation of AI's potential to impact the federal budget. While the nature and magnitude of these shocks vary, several implicit assumptions frame the collection of shocks. To start, we implicitly assume that AI's economic impact will be at least moderate—if not more substantial—and rising over time as the technology is adopted more widely and its capabilities continue to develop. In addition, we assume—in line with virtually every other major technological shock—that the net impact of widespread AI adoption will be productivity enhancing and lead to greater national incomes over time. However, we also note that the evidence to date suggests that AI may impact the fiscal outlook in substantially different channels than in prior technological revolutions. Specifically, while AI may ultimately have a profound impact on productivity, AI has already shown the potential for dramatically changing health care delivery, effectiveness, and cost—which could translate into changes in mortality, morbidity, price of care, and care utilization. Given that such changes could have profound impacts on Social Security and public health program outlays, policymakers would benefit from proactively integrating AI's capabilities into fiscal planning and projections.

In this paper, we model the potential impact of AI on Social Security outlays, Medicare outlays, individual and corporate tax bases, and the subsequent change in net interest payments. These components of the federal budget comprise the vast bulk of revenue and outlays, with an increasing share over time with the expansion of old-age entitlements as a share of the federal budget.[2] Because our analysis covers such a large share of fiscal policy, and we believe these fiscal elements are the most likely to be impacted by AI over the next two decades, we are comfortable characterizing the results as the impact on the fiscal outlook as a whole. However, there is potential for AI to impact various other fiscal elements, including payroll taxes to fund major entitlements, defense spending, other health programs such as Medicaid, premium support, and the Children's Health Insurance Program. Analyzing the impact of AI on these fiscal elements presents an opportunity for future research.

---

[2] Congressional Budget Office. 2024. "The Long-Term Budget Outlook: 2024 to 2054." https://www.cbo.gov/publication/59711.

Recent Developments and Medical Applications

The field of AI has evolved rapidly over the last decade. Particularly significant advances in AI technology include the 2017 development of the Transformer model.[3] The latter breakthrough provided the foundation for modern Large AI Models (LAMs) such as OpenAI's Generative Pre-Trained Transformer 4 (GPT-4).

LAMs (or colloquially large language models (LLMs)) based on Transformer have continuously progressed in size, complexity, and capability. In 2023, OpenAI released their GPT-4 model, which is rumored to have 1.76 trillion parameters,[4] making it hundreds of times larger than the 340 million parameters of Google's Bidirectional Encoder Representations from Transformers (BERT) LAM model released in 2018.[5] While model size is imperfectly correlated to capability, larger AI models tend to be more capable than smaller ones. AI models' number of "modalities" has also expanded; while earlier LAMs were only be trained on language data (LLMs) or image data (large vision models; LVMs), LAMs can now be trained on two or more types of data (large multi-modal models; LMMs).[6]

On the less technical front, AI products like ChatGPT have seized the public's imagination and taken AI mainstream as a topic of both excitement and concern. This widespread discussion has intensified existing academic interest in AI's applications in a variety of fields, including law,[7] finance,[8] economics,[9] and most significantly for our purposes, medicine.

AI products have already impacted health care. For example, the Alphafold 2 program, based on a Transformer model, has revolutionized protein structure prediction ("protein folding") since its release in 2020. Before Alphafold, decades of experimentation had left researchers with a complete structural understanding of only about 17 percent of the protein residues in the human body. In contrast, Alphafold was able to quickly develop a confident structural prediction for 58 percent of proteins.[10] Alphafold offers researchers unprecedented insight into the building blocks of human life, which could accelerate the speed of medical research and drug discovery.[11]

LAMs are also increasingly being leveraged towards improved medical diagnoses, an area in which they show incredible promise. Google's Articulate Medical Intelligence Explorer (AMIE), an LLM-based system, performed better than human clinicians when evaluating over 300 challenging diagnostic cases drawn from the New England Journal of Medicine. The LLM listed

---

[3] Ashish Vaswani et al., "Attention Is All You Need" (arXiv, August 1, 2023), https://doi.org/10.48550/arXiv.1706.03762.

[4] Schreiner, Maximilian, "GPT-4 architecture, datasets, costs and more leaked", *the decoder* (July 11, 2023), https://the-decoder.com/gpt-4-architecture-datasets-costs-and-more-leaked/; Stern, Jacob, "GPT-4 Might Just Be a Bloated Pointless Mess", *The Atlantic* (March 6, 2023), https://www.theatlantic.com/technology/archive/2023/03/openai-gpt-4-parameters-power-debate/673290/.

[5] Nvidia, "BERT". Accessed October 15, 2024. https://www.nvidia.com/en-us/glossary/bert/.

[6] Jianing Qiu et al., "Large AI Models in Health Informatics: Applications, Challenges, and the Future," *IEEE Journal of Biomedical and Health Informatics* 27, no. 12 (December 2023): 6074–87, https://doi.org/10.1109/JBHI.2023.3316750.

[7] John Armour and Mari Sako, "AI-Enabled Business Models in Legal Services: From Traditional Law Firms to next-Generation Law Companies?," *Journal of Professions and Organization* 7, no. 1 (March 1, 2020): 27–46, https://doi.org/10.1093/jpo/joaa001.

[8] John W. Goodell et al., "Artificial Intelligence and Machine Learning in Finance: Identifying Foundations, Themes, and Research Clusters from Bibliometric Analysis," *Journal of Behavioral and Experimental Finance* 32 (December 1, 2021): 100577, https://doi.org/10.1016/j.jbef.2021.100577.

[9] Anton Korinek, "Generative AI for Economic Research: Use Cases and Implications for Economists," *Journal of Economic Literature* 61, no. 4 (January 2023): 1281–1317, https://doi.org/10.1257/jel.20231736.

[10] Kathryn Tunyasuvunakool et al., "Highly Accurate Protein Structure Prediction for the Human Proteome," *Nature* 596, no. 7873 (August 2021): 590–96, https://doi.org/10.1038/s41586-021-03828-1.

[11] Ibid.

the correct diagnosis among its top-10 predictions 59.1 percent of the time, significantly outscoring human clinicians' 33.6 percent top-10 accuracy.[12] AMIE also excels at interacting with patients; patient actors scored text-based consultations with AMIE as being significantly better than those provided by primary care physicians across the vast majority of evaluation axes, including empathy and sensitivity.[13]

Beyond protein folding and diagnosis, current LAMs can perform better than previously state of the art methods in categories of health care tasks such as medical imaging, medical informatics, medical education, public health.[14] For their part, Generative Adversarial Networks (GANs) are used for synthetic medical image generation, augmenting otherwise limited sets of training data and thereby improving the performance of other neural networks.[15] Furthermore, GAN-synthesized data can help anonymize health care data, ameliorating AI-related privacy concerns.[16] Although the future capabilities of LAMs is difficult to forecast, the growth in capabilities of frontier AI models in the past two years makes clear the direction of travel. As the capabilities of AI models expand, so do their implications for society through a variety of channels including health care and the federal budget.

Potential Impact of AI on Health Care and Longevity

One of AI's largest potential impacts will be in accelerating the efficacy of preventive medicine. The use of AI in preventive care and early detection of diseases could lead to a reduction in morbidity rates, contributing to a healthier population that requires less medical intervention over time. AI algorithms have shown remarkable success in diagnosing diseases from images (such as radiology scans) and predicting patient outcomes based on historical health data. AI's ability to improve diagnostic accuracy can not only improve patient outcomes but also reduce wasteful spending on inappropriate treatments. These tools can assist clinicians in detecting conditions earlier and with greater precision, potentially enabling earlier interventions that extend longevity.

Additionally, AI shows significant promise in optimizing treatment plans. By rapidly analyzing massive amounts of data from a wide range of sources, AI can help identify the most effective and cost-efficient individualized treatment plans for patients. This includes determining which medications are likely to be most effective based on a patient's unique profile, thus avoiding costly and ineffective treatments.

Similarly, AI applications in monitoring patient health and predicting flare-ups of chronic conditions can lead to better management of chronic diseases and reduce the need for expensive hospitalizations and treatments. Wearable devices and mobile health apps, powered by AI, enable real-time monitoring and can alert patients and health care providers to potential health issues before they require more serious intervention. The aim is to integrate data from wearable devices, patient records, and call transcripts into unified systems that act as "co-pilots" for health care

[12] Daniel McDuff et al., "Towards Accurate Differential Diagnosis with Large Language Models" (arXiv, November 30, 2023), https://doi.org/10.48550/arXiv.2312.00164.

[13] Tao Tu et al., "Towards Conversational Diagnostic AI" (arXiv, January 10, 2024), https://doi.org/10.48550/arXiv.2401.05654.

[14] Jianing Qiu et al., "Large AI Models in Health Informatics: Applications, Challenges, and the Future," *IEEE Journal of Biomedical and Health Informatics* 27, no. 12 (December 2023): 6074–87, https://doi.org/10.1109/JBHI.2023.3316750.

[15] ChangHyuk Kwon et al., "Increasing Prediction Accuracy of Pathogenic Staging by Sample Augmentation with a GAN," *PLOS ONE* 16, no. 4 (April 27, 2021): e0250458, https://doi.org/10.1371/journal.pone.0250458.

[16] Esteban Piacentino, Alvaro Guarner, and Cecilio Angulo, "Generating Synthetic ECGs Using GANs for Anonymizing Healthcare Data," *Electronics* 10, no. 4 (January 2021): 389, https://doi.org/10.3390/electronics10040389.

providers, keeping them informed about their patients' conditions in real-time. By reducing the need for in-person health care, this can alleviate capacity constraints across the entire health care system.

These advances in AI have the potential to dramatically alter the scope of federal spending on old-age entitlement programs, which can subsequently alter the fiscal trajectory. From a more optimistic perspective, existing AI systems may lower expenditures on all health spending, including Medicare, with cost reductions occurring through several channels—with personalized medicine being a prominent example. AI enables the analysis of vast amounts of data, including genetic information, lifestyle factors, and environmental exposures, to tailor treatments to individual patients. This personalized approach may significantly improve outcomes by targeting therapies that are most likely to be effective for a particular patient, reducing the trial-and-error approach that characterizes much of current medical practice. AI may further reduce health care costs by avoiding unnecessary treatments and hospital admissions, thus lowering the financial burden on the public health care system. AI can also help identify and prevent fraudulent Medicare claims, saving costs for the program.

Beyond direct patient care, AI may enhance health care quality by improving hospital and clinic operations. From optimizing appointment scheduling to managing patient flow and predicting peak times for different services, AI can help reduce wait times and improve the patient experience. Similarly, AI could potentially automate administrative tasks such as data entry, appointment scheduling, and even preliminary data analysis for diagnostic purposes. By reducing the burden of repetitive tasks on health care professionals, AI allows them to focus more on patient care, thereby increasing the efficiency of health care delivery and reducing labor costs.

The rapid investments in AI-based health care by the world's leading technology companies signify a pivotal shift towards a more efficient, patient-centered approach. Moreover, the pursuit of health care innovation by these companies underscores a burgeoning competition that could significantly enhance medical care and operational efficiency. For example, Google, a leader in leveraging AI for health advancements, is at the forefront with projects like Med-Gemini. This health-specific LLM focuses on streamlining health care by providing accurate responses to medical inquiries and facilitating the summarization of vital information during pivotal moments such as patient handoffs and staff shift changes. Microsoft's recent string of strategic acquisitions further exemplifies the tech industry's drive towards enhancing health care through AI. In 2021, Microsoft acquired Nuance, which is designed to assist health care professionals with administrative tasks such as generating clinical notes and managing electronic health records. Similarly, Amazon's collaboration with Anthropic aims to introduce a version of Claude to augment health care services, showcasing Amazon's commitment to enriching health care delivery through technological empowerment.

The race to harness AI for health care advantages isn't just happening among tech giants—it's also a point of competition and collaboration among nations; Chinese technology giants are also venturing into this arena. A 2022 McKinsey report highlighted AI's potential to revolutionize health care in China, projecting that AI's integration into diagnostic predictions and clinical decision support could generate approximately $5 billion in economic value.[17] This underscores a global recognition of AI's transformative potential in health care.

The effects of improved public health through AI-driven health care initiatives extend far beyond the immediate benefits of reduced disease burden and decreased health care costs. By

[17] Shen, Kai, Xiaoxiao Tong, Ting Wu, and Fangning Zhang. "The next frontier for AI in China could add $600 billion to its economy." *QuantumBlack AI by McKinsey* (June 7, 2022). https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-next-frontier-for-ai-in-china-could-add-600-billion-to-its-economy.

significantly lowering morbidity and mortality rates, AI has the potential to bolster labor supply, including both the supply of working-age individuals who may experience fewer sick days, but also the potential to dramatically increase labor force participation by older households who may experience increased health in old age and expanded longevity.

These positive developments are not guaranteed to translate into fiscal improvements. Under a scenario where AI leads to reductions in mortality rates, but not cost savings in per beneficiary health costs, the fiscal outlook could deteriorate as an expanded old-age population implies higher fiscal outlays. Moreover, improved efficacy of health care delivery could potentially increase health care utilization by driving up demand for such services—although such a scenario would likely be accompanied by price reductions due to improved health care productivity. In short, the AI revolution on health care could counterintuitively increase both per capita spending on entitlements and the population of beneficiaries receiving these services.

AI could, of course, impact the economy in more traditional ways, including in particular by bolstering labor and total factor productivity—similar to the impact of information technology in the late 1990s. This more traditional shock could drive up tax revenues owing to higher real incomes and profits, thus relaxing fiscal pressure, all else equal.[18]

This working paper is based on exploratory analysis, which attempts to place structure on forecasts of AI's potential impact on medicine and fiscal budgets. Our analyses should be regarded as an initial attempt to scope the potential magnitude of an AI shock on the long-term federal budget outlook—which we consider to be a 20-year time frame. Our initial estimates suggest that the nature of the shock is critical, as the impact of the shock on annual budget deficits could range from an increase of 0.9 percent of gross domestic product (GDP) to a decrease of 3.8 percent of GDP, with the latter instance effectively halving annual budget deficits. In the next section we review the literature around the impact of AI on various aspects of health care and longevity. Section 3 presents a theoretical model. In Section 4 we lay out our methodology, and our results are presented in Section 5. Section 6 briefly concludes.

2. Literature Review

AI can potentially impact the budget outlook through three health care-related channels: health care utilization, health care costs, and longevity. AI is also likely to have more general, non-health care related impacts on productivity, which will have their own fiscal effects. In the following section we examine the literature supporting each channel. We also examine comparable historical shocks to conceptualize a reasonable range for the magnitude of the impact AI may have within each channel.

Health Care Utilization

The use of AI presents the rare—possibly unique—opportunity to expand access to health care information and services while simultaneously reducing the burden on the conventional health care system. AI tools can accomplish this by making diagnoses faster and more accurate, facilitating the prioritization of higher-risk patients, and defending against future public health

[18] AI could potentially lead to labor market disruptions, which may put significant new strains on social safety nets and further deteriorate the fiscal outlook (e.g., Klinova, Katya, and Anton Korinek. "Unleashing Possibilities, Ignoring Risks: Why We Need Tools to Manage AI's Impact on Jobs." Brookings Institution, August 17, 2023. https://www.brookings.edu/articles/unleashing-possibilities-ignoring-risks-why-we-need-tools-to-manage-ais-impact-on-jobs/.).

crises—all while empowering people without medical expertise to have more agency over their individual health outcomes.

Although there is an extensive literature surrounding the use of AI tools in various fields of medicine, few studies attempt to estimate the extent to which AI will impact overall health care utilization. Consequently, in the following subsection we summarize studies that examine specific cases of AI technologies that have the potential to affect health care utilization. These narrow applications, if taken together, allow us to characterize the broader contours of the AI shock to health care utilization.

AI will facilitate more accurate medical screening and earlier diagnosis. This could allow health care providers to head-off problems that, if identified later, would require more dramatic interventions or could develop into chronic illnesses. A notable example of an illness that AI can address is diabetic retinopathy (DR), the leading cause of blindness in working-age adults worldwide.[19] Abràmoff et al. (2018) demonstrated that an AI screening system for DR, "IDx-DR," achieved a diagnostic sensitivity that exceeded the average sensitivity of board-certified ophthalmologists. Backed by these results, IDx-DR became the first ever fully autonomous system to be approved by the Food and Drug Administration for use in medicine.[20]

AI-based screening systems can assist clinicians in earlier identification of several other diseases, as well as highlighting high-risk patients for closer diagnostic attention and treatment. Hill et al. (2022) developed an AI model capable of identifying patients with an elevated risk of atrial fibrillation, a common form of arrhythmia which affects millions of Americans and is associated with an increased likelihood of stroke, heart failure, and premature cognitive decline.[21,22] Issaiy, Zarei, and Saghazadeh (2023) conducted a systematic review of the use of AI algorithms in the diagnosis and prognosis of acute appendicitis, finding that AI algorithms achieved high sensitivity and specificity, and often surpassed the speed and accuracy of traditional diagnostic methods.[23]

The field of radiology is particularly receptive to the application of AI. Tools leveraging AI can accelerate patient evaluations or reduce the burden on clinicians when identifying injuries and illnesses. Kim et al. (2020) developed an AI algorithm that demonstrated superior performance in breast cancer diagnosis compared to human radiologists.[24] Annarumma et al. (2019) showed that AI assistance reduced the time it took for patients to receive radiologists' interpretations of chest X-

---

[19] Cheung, Ning, Paul Mitchell, and Tien Yin Wong. "Diabetic Retinopathy." *The Lancet* 376, no. 9735 (July 10, 2010): 124–36. https://doi.org/10.1016/S0140-6736(09)62124-3.

[20] Abràmoff, Michael D., Philip T. Lavin, Michele Birch, Nilay Shah, and James C. Folk. 2018. "Pivotal Trial of an Autonomous AI-Based Diagnostic System for Detection of Diabetic Retinopathy in Primary Care Offices." *Npj Digital Medicine* 1 (1): 1–8. https://doi.org/10.1038/s41746-018-0040-6.

[21] Hill, Nathan R, Lara Groves, Carissa Dickerson, Andreas Ochs, Dong Pang, Sarah Lawton, Michael Hurst, et al. 2022. "Identification of Undiagnosed Atrial Fibrillation Using a Machine Learning Risk-Prediction Algorithm and Diagnostic Testing (PULsE-AI) in Primary Care: A Multi-Centre Randomized Controlled Trial in England." *European Heart Journal - Digital Health* 3 (2): 195–204. https://doi.org/10.1093/ehjdh/ztac009.

[22] Pipilas, Daniel, Samuel Freesun Friedman, and Shaan Khurshid. 2023. "The Use of Artificial Intelligence to Predict the Development of Atrial Fibrillation." *Current Cardiology Reports* 25 (5): 381–89. https://doi.org/10.1007/s11886-023-01859-w.

[23] Issaiy, Mahbod, Diana Zarei, and Amene Saghazadeh. 2023. "Artificial Intelligence and Acute Appendicitis: A Systematic Review of Diagnostic and Prognostic Models." *World Journal of Emergency Surgery* 18 (1): 59. https://doi.org/10.1186/s13017-023-00527-2.

[24] Kim, Hyo-Eun, Hak Hee Kim, Boo-Kyung Han, Ki Hwan Kim, Kyunghwa Han, Hyeonseob Nam, Eun Hye Lee, and Eun-Kyung Kim. "Changes in Cancer Detection and False-Positive Recall in Mammography Using Artificial Intelligence: A Retrospective, Multireader Study." *The Lancet Digital Health* 2, no. 3 (March 1, 2020): e138–48. https://doi.org/10.1016/S2589-7500(20)30003-0.

rays from roughly 11 days to only about three days.[25] Finally, Arbabshirani et al. (2018) trained a neural network that reduced the time required to diagnose patients with intracranial hemorrhaging by 96 percent by rapidly analyzing computerized tomography (CT) scans.[26]

The speed and accuracy of AI-based (and AI-assisted) screening and diagnosis systems could reduce diagnostic errors, which are a massive cause of serious injury in the United States. Newman-Toker et al. (2024) estimated that each year diagnostic error is responsible for 795,000 serious harms in the United States, among which are 371,000 deaths and 424,000 permanent disabilities.[27] These hundreds of thousands of permanently disabled people will require additional medical care for the remainder of their lives; Khavjou et al. (2020) concluded that, in 2015, per capita medical expenditures on people with disabilities were 2.5 times greater than those without disabilities.[28]

Cutting down on injuries caused by late or mis- diagnoses is just one of the avenues through which health care utilization would fall under the new paradigm of AI-driven medicine. Hospital readmissions could also be driven down by machine-learning models that facilitate the prioritization of higher-risk patients. Mohanty et al. (2021) developed an AI model that could predict patients' risk of readmission by considering parameters such as patients' demographic characteristics, comorbidities, and level of frailty.[29] These efforts are especially vital for older patients, for whom readmission is an especially serious issue.

AI could further reduce strain on the American health care system by assisting in the prevention and management of future public health events such as pandemics. Brownstein et al. (2023) and Olawade et al. (2023) review the use of AI in public health, documenting applications in the forecasting and spatial-modeling of disease outbreaks, as well as public health surveillance, misinformation control, and the allocation of limited testing resources.[30,31] Gadaleta et al. (2021) demonstrated that a machine learning algorithm could even identify presymptomatic COVID-19 infections using subclinical changes documented by patients' smartwatches, suggesting that AI might allow the recognition of disease hotspots before more serious symptoms appear.[32]

---

[25] Annarumma, Mauro, Samuel J. Withey, Robert J. Bakewell, Emanuele Pesce, Vicky Goh, and Giovanni Montana. 2019. "Automated Triaging of Adult Chest Radiographs with Deep Artificial Neural Networks." *Radiology* 291 (1): 196–202. https://doi.org/10.1148/radiol.2018180921.

[26] Arbabshirani, Mohammad R., Brandon K. Fornwalt, Gino J. Mongelluzzo, Jonathan D. Suever, Brandon D. Geise, Aalpen A. Patel, and Gregory J. Moore. 2018. "Advanced Machine Learning in Action: Identification of Intracranial Hemorrhage on Computed Tomography Scans of the Head with Clinical Workflow Integration." *Npj Digital Medicine* 1 (1): 1–7. https://doi.org/10.1038/s41746-017-0015-z.

[27] Newman-Toker, David E, Najlla Nassery, Adam C Schaffer, Chihwen Winnie Yu-Moe, Gwendolyn D Clemens, Zheyu Wang, Yuxin Zhu, et al. 2024. "Burden of Serious Harms from Diagnostic Error in the USA." *BMJ Quality & Safety* 33 (2): 109–20. https://doi.org/10.1136/bmjqs-2021-014130.

[28] Khavjou, Olga A., Wayne L. Anderson, Amanda A. Honeycutt, Laurel G. Bates, Hilda Razzaghi, NaTasha D. Hollis, and Scott D. Grosse. 2020. "National Health Care Expenditures Associated With Disability." *Medical Care* 58 (9): 826–32. https://doi.org/10.1097/MLR.0000000000001371.

[29] Mohanty, Somya D., Deborah Lekan, Thomas P. McCoy, Marjorie Jenkins, and Prashanti Manda. 2021. "Machine Learning for Predicting Readmission Risk among the Frail: Explainable AI for Healthcare." *Patterns* 3 (1): 100395. https://doi.org/10.1016/j.patter.2021.100395.

[30] Brownstein, John S., Benjamin Rader, Christina M. Astley, and Huaiyu Tian. 2023. "Advances in Artificial Intelligence for Infectious-Disease Surveillance." *New England Journal of Medicine* 388 (17): 1597–1607. https://doi.org/10.1056/NEJMra2119215.

[31] Olawade, David B., Ojima J. Wada, Aanuoluwapo Clement David-Olawade, Edward Kunonga, Olawale Abaire, and Jonathan Ling. "Using Artificial Intelligence to Improve Public Health: A Narrative Review." *Frontiers in Public Health* 11 (October 26, 2023): 1196397. https://doi.org/10.3389/fpubh.2023.1196397.

[32] Gadaleta, Matteo, Jennifer M. Radin, Katie Baca-Motes, Edward Ramos, Vik Kheterpal, Eric J. Topol, Steven R. Steinhubl, and Giorgio Quer. 2021. "Passive Detection of COVID-19 with Wearable Sensors and Explainable Machine Learning Algorithms." *Npj Digital Medicine* 4 (1): 1–10. https://doi.org/10.1038/s41746-021-00533-1.

The utilization of the traditional health care system would naturally fall if people could safely assess and manage their own health. AI tools may one day permit this, meaning that fewer people would depend on the expertise of medical professionals. AI-based smartphone applications (apps) have already proven capable of assisting in the non-invasive screening of many diseases. The authors of Mannino et al. (2018) created an app that could diagnose anemia using only photos of patients' fingernail beds.[33] Singh and Xu (2024) developed another app capable of diagnosing whether a patient had Parkinson's disease with 99 percent accuracy in under a second using only a ten second audio recording.[34]

Besides diagnostic apps, AI can empower people to manage their own health by providing other health services. Chew (2022) reviewed 23 studies featuring AI chatbots used for weight loss.[35] While only four of the included studies reported on the effectiveness of chatbots on user outcomes like diet, physical activity, or weight loss, three of those studies reported improved outcomes in the chatbot programs compared to control groups. Despite these limited findings, the author suggests that in the coming years, AI providing exercise and nutrition recommendations have the potential to cut down on the rates of obesity. AI's ability to alter patient behavior might also improve the effectiveness of prescriptions without having to change the formulation or dosage of the drug itself. Ilan (2021) proposes coupling drugs with a personalized AI system, creating a "digital pill" that will improve adherence to drug regimens.[36]

AI's most impactful change may come through the personalization of medicine. Musich et al. (2016) compared thousands of patients who were members in a network of affiliated primary care physicians focused on personalized preventative health care to a matched set of patients who were not members.[37] The network of physicians provided services which included detailed health screenings, diagnostics, and personalized nutrition and exercise coaching—the exact services we see AI beginning to provide. Prevention-focused health management programs' impact on the utilization of emergency room and urgent care services might therefore act as an analog for the overall impact of AI on health care utilization. Musich et al. (2016) also found that after three years, network members were statistically significantly less likely to have had an emergency room visit, as well as being less likely to have used an urgent care facility.[38]

Expanded access to screening and improved diagnostic accuracy can catch illnesses early and prevent the development of more dangerous, difficult-to-treat symptoms. However, more frequent diagnosis could also drive increased health care utilization. There is growing concern about "overdiagnosis," which occurs when an asymptomatic illness is diagnosed despite not

---

[33] Mannino, Robert G., David R. Myers, Erika A. Tyburski, Christina Caruso, Jeanne Boudreaux, Traci Leong, G. D. Clifford, and Wilbur A. Lam. 2018. "Smartphone App for Non-Invasive Detection of Anemia Using Only Patient-Sourced Photos." *Nature Communications* 9 (1): 4924. https://doi.org/10.1038/s41467-018-07262-2.

[34] Singh, Sanjana, and Wenyao Xu. 2020. "Robust Detection of Parkinson's Disease Using Harvested Smartphone Voice Data: A Telemedicine Approach." *Telemedicine Journal and E-Health* 26 (3): 327. https://doi.org/10.1089/tmj.2018.0271.

[35] Chew, Han Shi Jocelyn. 2022. "The Use of Artificial Intelligence–Based Conversational Agents (Chatbots) for Weight Loss: Scoping Review and Practical Recommendations." *JMIR Medical Informatics* 10 (4): e32578. https://doi.org/10.2196/32578.

[36] Ilan, Yaron. 2021. "Improving Global Healthcare and Reducing Costs Using Second-Generation Artificial Intelligence-Based Digital Pills: A Market Disruptor." *International Journal of Environmental Research and Public Health* 18 (2): 811. https://doi.org/10.3390/ijerph18020811.

[37] Musich, Shirley, Shaohung Wang, Kevin Hawkins, and Andrea Klemes. 2016. "The Impact of Personalized Preventive Care on Health Care Quality, Utilization, and Expenditures." *Population Health Management* 19 (6): 389–97. https://doi.org/10.1089/pop.2015.0171.

[38] In particular, the study found that while 17.3 percent of nonmembers visited the emergency room (ER), only 14.6 percent of members had to visit the ER. Even more extreme, while 10.1 percent of nonmembers used an urgent care facility, only 4.5 percent of members did the same.

causing a patient any pain or discomfort. Dunn et al. (2022) reviewed overdiagnosis in cancer screening, noting that overdiagnosis can lead to physical and financial harm from unnecessary treatments, as well as the psychological trauma of being labeled as sick. Furthermore, overdiagnosis contributes to "overmedicalization," wherein normal life experiences are treated as symptoms of disease.[39] Overdiagnosis already occurs frequently. Bleyer and Welch (2012) concluded that 1.3 million American women had been overdiagnosed with breast cancer over the preceding 30 years.[40]

Health Care Costs

Many studies have examined the cost-effectiveness of various applications of AI in health care, particularly in their diagnostic role. Returning to the example of diabetic retinopathy (DR), Fuller et al. (2022) found that an automated DR screening system reduced costs by 23.3 percent compared to the current standard of care screening systems, while performing just as well.[41] Similarly, Wolf et al. (2020) modeled that autonomous DR screening methods would generate cost-savings so long as more than 23 percent of diabetic patients adhered to yearly screening recommendations—though AI screening was more costly if adherence was below 23 percent.[42]

As for other diseases, a model developed by Pickhardt, Correale, and Hassan (2023) demonstrated that AI assistance was cost-saving when applied to CT-based screenings for cardiovascular diseases, osteoporosis, and sarcopenia.[43] Areia et al. (2022) employed a microsimulation model to compare the costs of screening colonoscopies with and without AI assistance. Within the model, AI-assisted screening prevented 7,194 cases of colorectal cancer, prevented 2,089 deaths, and saved $290 million annually.[44]

Although most of the applications of AI that we have examined thus far have been directly related to patient health, one of the most important ways in which AI can reduce health care prices is by cutting administrative costs. Sahni, Carrus, and Cutler (2021) estimated that one-quarter of the approximately $3.8 trillion the United States spent on health care in 2019 went towards administrative functions.[45] Sahni et al. (2023) calculated that reduced administrative costs would

[39] Dunn, Barbara K., Steven Woloshin, Heng Xie, and Barnett S. Kramer. "Cancer Overdiagnosis: A Challenge in the Era of Screening." *Journal of the National Cancer Center* 2, no. 4 (December 1, 2022): 235–42. https://doi.org/10.1016/j.jncc.2022.08.005.

[40] Bleyer Archie and Welch H. Gilbert. "Effect of Three Decades of Screening Mammography on Breast-Cancer Incidence." *New England Journal of Medicine* 367, no. 21 (2012): 1998–2005. https://doi.org/10.1056/NEJMoa1206809.

[41] Fuller, Spencer D., Jenny Hu, James C. Liu, Ella Gibson, Martin Gregory, Jessica Kuo, and Rithwick Rajagopal. "Five-Year Cost-Effectiveness Modeling of Primary Care-Based, Nonmydriatic Automated Retinal Image Analysis Screening Among Low-Income Patients With Diabetes." *Journal of Diabetes Science and Technology* 16, no. 2 (March 1, 2022): 415–27. https://doi.org/10.1177/1932296820967011.

[42] Wolf, Risa M., Roomasa Channa, Michael D. Abramoff, and Harold P. Lehmann. "Cost-Effectiveness of Autonomous Point-of-Care Diabetic Retinopathy Screening for Pediatric Patients With Diabetes." *JAMA Ophthalmology* 138, no. 10 (October 1, 2020): 1063–69. https://doi.org/10.1001/jamaophthalmol.2020.3190.

[43] Pickhardt, Perry J., Loredana Correale, and Cesare Hassan. "AI-Based Opportunistic CT Screening of Incidental Cardiovascular Disease, Osteoporosis, and Sarcopenia: Cost-Effectiveness Analysis." *Abdominal Radiology* 48, no. 3 (March 1, 2023): 1181–98. https://doi.org/10.1007/s00261-023-03800-9.

[44] Areia, Miguel, Yuichi Mori, Loredana Correale, Alessandro Repici, Michael Bretthauer, Prateek Sharma, Filipe Taveira, et al. "Cost-Effectiveness of Artificial Intelligence for Screening Colonoscopy: A Modelling Study." *The Lancet Digital Health* 4, no. 6 (June 1, 2022): e436–44. https://doi.org/10.1016/S2589-7500(22)00042-5.

[45] Sahni, Nikhil R., Brandon Carrus, and David M. Cutler. "Administrative Simplification and the Potential for Saving a Quarter-Trillion Dollars in Health Care." *JAMA* 326, no. 17 (November 2, 2021): 1677–78. https://doi.org/10.1001/jama.2021.17315.

account for more than a third of the $200-$360 billion (in 2019 dollars) that the authors predicted AI would cut from health care spending over the five years following that paper's publication.[46]

Despite mounting evidence of the cost-effectiveness of AI in health care, the novelty of AI technology means that there has been limited empirical study of the impact of AI on the price of health services for consumers. As a result, it remains unclear to what degree, if any, the cost savings associated with AI will be passed on to patients and payers, and to what degree they will be internalized by health care providers. That being said, studies have examined the impact of policy reforms on the trajectory of health care prices, as well as the extent to which spending has changed due to medical and technical advancements.

On the policy front, the Affordable Care Act (ACA) is a historical example of a shock that substantially altered the trajectory of health care prices. The ACA aimed to make health insurance more widely available, expand the coverage of Medicaid, and lower the general cost of health care through innovative medical care delivery methods.[47] The ACA has had several consequences to the price of health care in the United States, including reducing the growth rate of cumulative Medicare prices. According to Buntin and Graves (2020), while the overall index of Medicare prices increased by 14.2 percent between 2010 and 2018, this increase would have been 5.2 percentage points higher (19.4 percent) without reductions from the ACA.[48]

The benefits of technological advancement on the price of care are not firmly established. In fact, there is a substantial literature which suggests that technological advancement is a major cause of increasing health care costs. Conducting a literature review of studies on the impact of technology on health expenditure growth, Marino and Lorenzoni (2019) calculated that the literature attributed an average of 35 percent of the growth of health expenditures to technological change.[49] The authors explained that, while many new technologies reduce costs per treatment, breakthroughs also increase spending by rendering existing treatments irrelevant, treating previously untreatable diseases, and expanding the use of certain treatments through more accurate and accessible diagnoses.

And yet, AI tools differ from previous technological transformations. Typically, medical advancements have only affected which ailments were treatable, and how they could be treated; the responsibility of administering care has always remained in the hands of the health care professionals. In contrast, AI tools present an opportunity for the "democratization" of health care; changing the "who" and "where" of preventative medical care. As previously discussed, enabling consumers to effectively manage their personal health may decrease the utilization of the formal medical system. In addition to a straightforward reduction in spending on services such as checkups and basic consultations, the reduced demand for health care caused by AI tools has to the potential to shift the demand curve for care inwards, further driving down care prices for all consumers. Alternatively, technological advances may increase efficacy and thus utilization, potentially raising per capita expenditures in the process.

Death Rates and Longevity

---

[46] Sahni, Nikhil, George Stein, Rodney Zemmel, and David M. Cutler. "The Potential Impact of Artificial Intelligence on Healthcare Spending." Working Paper. Working Paper Series. National Bureau of Economic Research, January 2023. https://doi.org/10.3386/w30857.

[47] "Affordable Care Act (ACA) - Glossary." n.d. HealthCare.Gov. Accessed April 23, 2024. https://www.healthcare.gov/glossary/affordable-care-act.

[48] Buntin, Melinda Beeuwkes, and John A. Graves. 2020. "How The ACA Dented The Cost Curve." *Health Affairs* 39 (3): 403–12. https://doi.org/10.1377/hlthaff.2019.01478.

[49] Marino, Alberto, and Luca Lorenzoni. "The Impact of Technological Advancements on Health Spending: A Literature Review." Paris: OECD, August 22, 2019. https://doi.org/10.1787/fa3bab05-en.

The evidence around expanding access to health care and lowered health care prices, coupled with the promise of widespread application of personalized medicine and more effective diagnostic and treatment procedures, suggests that AI may have a marked impact on longevity—although there is not enough evidence to make a decisive conclusion on the magnitude of the impact.

Instead, we look to history. In the past century there have been instances of extremely rapid declines in mortality. A salient example occurred in Japan, where the increase in life expectancy was especially dramatic in the years following the Second World War (WWII). The average life expectancy at birth of Japanese citizens grew by about 13.7 years in the years immediately following the war (1947-55).[50]

Rapid growth in life expectancies and declines in mortality are driven by several factors. In Japan's case, Sugiura et al. (2007) attributed the rise in life expectancy following WWII to a greater intake of protein, improved health education and regular physical checkups in schools, expanded health laws and regulations (particularly regarding qualified medical staff), and drastic agricultural reforms.[51] Since then, Japan has become the longest-lived major economy, which Tsugane (2021) argued is due to very low mortality rates from ischemic heart disease, breast cancer, and prostate cancer. Tsugane (2021) noted that these low mortality rates are believed to be consequences of the low rate of obesity, which is itself a consequence of diet.[52] The AI weight-loss chatbots discussed in the previous subsection may therefore cause increased longevity if they can successfully encourage users to improve their diets.

A positive shock to longevity, all else equal, can markedly increase federal spending. While extended longevity in isolation will unambiguously increase federal outlays, the impact on the federal budget also depends on assumptions about health care utilization and prices and labor supply decisions. Dieleman et al. (2017) calculated that population aging was associated with an 11.6 percent increase in health care spending between 1996 and 2013, behind population growth (23.1 percent) and rising price and intensity of services (50.0 percent).[53] Even as the baby-boomer generation ages into retirement, Keehan et al. (2017) projected that the American population's changing age-sex mix would contribute less to health care spending growth from 2020 to 2025 than rising medical prices, the increasing use and intensity of medical services, and population growth.[54]

Non-health Productivity

The widespread adoption of AI tools can potentially have fiscal impacts outside health care. For example, increased productivity could lead to increased wage and capital income, which could expand the revenue collected through the individual income and corporate tax systems. While there

---

[50] "Abridged Life Tables for Japan 2019." Director-General for Statistics and Information Policy, Ministry of Health, Labour and Welfare. Government of Japan. Accessed April 23, 2024. https://www.mhlw.go.jp/english/database/db-hw/lifetb19/dl/lifetb19-06.pdf.

[51] Sugiura, Yasuo, Young-Su Ju, Junko Yasuoka, and Masamine Jimba. "Rapid Increase in Japanese Life Expectancy after World War II." *Bioscience Trends* 4, no. 1 (February 2010): 9–16.

[52] Tsugane, Shoichiro. "Why Has Japan Become the World's Most Long-Lived Country: Insights from a Food and Nutrition Perspective." *European Journal of Clinical Nutrition* 75, no. 6 (June 2021): 921–28. https://doi.org/10.1038/s41430-020-0677-5.

[53] Dieleman, Joseph L., Ellen Squires, Anthony L. Bui, Madeline Campbell, Abigail Chapin, Hannah Hamavid, Cody Horst, et al. "Factors Associated With Increases in US Health Care Spending, 1996-2013." *JAMA* 318, no. 17 (November 7, 2017): 1668–78. https://doi.org/10.1001/jama.2017.15927.

[54] Keehan, Sean P., Devin A. Stone, John A. Poisal, Gigi A. Cuckler, Andrea M. Sisko, Sheila D. Smith, Andrew J. Madison, Christian J. Wolfe, and Joseph M. Lizonitz. "National Health Expenditure Projections, 2016–25: Price Increases, Aging Push Sector To 20 Percent Of Economy." *Health Affairs* 36, no. 3 (March 2017): 553–63. https://doi.org/10.1377/hlthaff.2016.1627.

is a great deal of interest in AI's potential effects on economic productivity, the magnitude of this impact remains unclear. Bailey, Brynjolfsson and Korinek (2023) draw on academic studies to articulate the case for an AI induced boom in productivity growth, seeing an 18 percent higher level of productivity over 10 years.[55] By contrast, Acemoglu (2024) finds far more limited gains from AI with a 0.7 percent increase in the level of productivity over 10 years.[56] Additionally, AI could potentially induce large shifts in the distribution of income between capital and labor that may partially or fully offset the productivity benefits for wage earners (see, for example, Bell and Korinek (2023) and Acemoglu (2023)).[57,58]

Estimates also abound from nonacademic sources. A McKinsey study authored by Chui et al. (2018) examined 400 use cases across 19 industries, estimating that AI had the potential to annually add \$3.5 trillion to \$5.8 trillion in value to those industries.[59] A Goldman Sachs report by Hatzius et al. (2023) concluded that, although generative AI has the potential to expose 300 million workers around the world to automation, it could also accelerate annual US labor productivity growth by around 1.5 percentage points over the next decade. Furthermore, widespread adoption of AI tools could increase annual global GDP by 7 percent.[60] Cazzaniga et al. (2024) conducted a model-based analysis of AI's economic impact, concluding that, while service-driven advanced economies are more exposed to AI than are developing economies, advanced economies are simultaneously better equipped to benefit from these new technologies.[61]

Recent studies have presented empirical evidence of AI improving productivity in specific settings. For example, Brynjolfsson, Li, and Raymond (2023) found that customer support agents were able to resolve 14 percent more issues per hour on average using an AI-based conversational assistant, with larger gains on for less skilled or experienced workers.[62] Another study, Noy and Zhang (2023), found that ChatGPT boosted productivity among college-educated professionals completing writing tasks, decreasing the time taken by 40 percent and improving quality by 18 percent.[63]

Historical precedent can also help scope the potential magnitude of AI's impact. For example, the 1990s saw a productivity shock brought on by the adoption of general-purpose

---

[55] Baily, Martin Neil, Erik Brynjolfsson, and Anton Korinek. "Machines of Mind: The Case for an AI-Powered Productivity Boom." Brookings, May 10, 2023. https://www.brookings.edu/articles/machines-of-mind-the-case-for-an-ai-powered-productivity-boom/.

[56] Acemoglu, Daron. "The Simple Macroeconomics of AI." Working Paper. Working Paper Series. National Bureau of Economic Research, May 2024. https://doi.org/10.3386/w32487.

[57] Bell, Stephanie A., and Anton Korinek. "AI's Economic Peril." *Journal of Democracy* 34, no. 4 (2023): 151–61. https://doi.org/10.1353/jod.2023.a907696.

[58] Acemoglu, Daron. "Harms of AI." In *The Oxford Handbook of AI Governance*, edited by Justin B. Bullock, Yu-Che Chen, Johannes Himmelreich, Valerie M. Hudson, Anton Korinek, Matthew M. Young, and Baobao Zhang, 660–706. Oxford Handbooks. Oxford University Press, 2023. https://doi.org/10.1093/oxfordhb/9780197579329.013.65.

[59] Chui, Michael, James Manyika, Mehdi Miremadi, Nicolaus Henke, Rita Chung, Pieter Nel, and Sankalp Malhotra. "Sizing the Potential Value of AI and Advanced Analytics," April 17, 2018. https://www.mckinsey.com/featured-insights/artificial-intelligence/notes-from-the-ai-frontier-applications-and-value-of-deep-learning.

[60] Hatzius, Jan, Joseph Briggs, Devesh Kodnani, and Giovanni Pierdomenico. "The Potentially Large Effects of Artificial Intelligence on Economic Growth." Goldman Sachs Economic Research, 2023. https://www.gspublishing.com/content/research/en/reports/2023/03/27/d64e052b-0f6e-45d7-967b-d7be35fabd16.html.

[61] Cazzaniga, Mauro, Florence Jaumotte, Longji Li, Giovanni Melina, Augustus J. Panton, Carlo Pizzinelli, Emma Rockall, and Marina M. Tavares. "Gen-AI: Artificial Intelligence and the Future of Work." International Monetary Fund, 2024. https://www.imf.org/en/Publications/Staff-Discussion-Notes/Issues/2024/01/14/Gen-AI-Artificial-Intelligence-and-the-Future-of-Work-542379.

[62] Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond. 2023. "Generative AI at Work." Working Paper. Working Paper Series. National Bureau of Economic Research. https://doi.org/10.3386/w31161.

[63] Noy, Shakked, and Whitney Zhang. "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence." *Science* 381, no. 6654 (July 14, 2023): 187–92. https://doi.org/10.1126/science.adh2586.

information technologies such as personal computers. Per capita labor productivity rose by an average of 9.6 percent per year from 1990 to 1995 and a whopping 13.1 percent per year from 1995 through 1999.[64]

[64] "High-Tech Productivity Gains in 1990s : The Economics Daily : U.S. Bureau of Labor Statistics." n.d. Bureau of Labor Statistics. Accessed April 23, 2024. https://www.bls.gov/opub/ted/2002/may/wk2/art02.htm.

3. Theoretical Framework

In this section, we layout a fiscal and macroeconomic accounting framework to assess the effects of generative AI on entitlement spending, including for major U.S. health care programs such as Medicare. This section describes a theoretical framework that is more general than our simulation but emphasizes the same basic channels—effects on mortality, utilization, price, and revenues through faster productivity growth. This framework can then be applied with more age-, industry-, or program-specific estimates relative to our more aggregated approach.

The first three channels reflect effects specific to health care: 1) AI may raise longevity, increasing the population eligible for Social Security and Medicare benefits, 2) AI could reduce morbidity implying lower health care utilization at any given age, but AI could also increase utilization through improved efficacy of health care delivery, 3) AI could lower the cost of health care by raising industry productivity or reducing industry labor costs, and 4) AI may have substantial revenue impacts through increases in capital or labor productivity that would raise tax revenue and potentially shift the distribution of income between capital and labor. In what follows, we make no assumptions on optimizing behavior by firms, workers, or households. We also do not impose any assumption on equilibrium prices needed to clear goods or capital markets.

Entitlement Expenditures

Let $N_{j,t}$ be the population of age j at time t. Then total population is given by $N_t = \sum_j N_{j,t}$. Let the survival rate at each age be given by $q_{j,t}$. The law of motion for each successive generation is given by $N_{j+1,t+1} = q_{j,t} N_{j,t}$. The retirement eligible population is given by $N_{ret,t} = \sum_{j \geq 65} N_{j,t} < N_t$.

Let Social Security expenditures for beneficiaries age j at time t be given by $SS_{j,t}$. Without loss of generality, Social Security expenditures per beneficiary at a given age can be defined as $ss_{j,t} = SS_{j,t}/N_{j,t}$. Then total old-age Social Security benefits are given by the expression below:

$$SS_{j,t} = ss_{j,t} N_{j,t}$$
$$\Rightarrow SS_t = \sum_{j \geq 65} SS_{j,t}$$

The theoretical framework assumes that effective health care expenditure at each age is given by $H_{j,t}$ at price $p_{m,t}$, where the subscript m refers to the medical or health care sector. The price of health services may vary over time, but a certain percentage of these expenditures x are covered by Medicare with the remainder covered by out-of-pocket expenses and health care premiums paid by beneficiaries.

Health care expenditures at each age and total retirement-age health care expenditures are given by the following expressions:

$$M_{j,t} = p_{m,t} H_{j,t} N_{j,t}$$
$$\Rightarrow M_t = \sum_{j \geq 65} M_{j,t}$$

Medicare expenditures are then simply a percentage of total health care expenditures.

To trace out the impact of generative AI on Medicare expenditures, we consider three potential channels: 1) effects on survival probabilities $q_{j,t}$, 2) effects on morbidity (i.e., reduction in health care services demand $H_{j,t}$, and 3) effects on the price of health care services $p_{m,t}$. As noted above, generative AI could raise health care costs by increasing survival probabilities and, thereby, increasing future Medicare and Social Security expenditures as beneficiaries live longer. Simultaneously, generative AI can lower the price of health care services over time by either improving the delivery of health care services or lowering labor costs. At the same time, generative AI may operate primarily by reducing demand for health care services by improving well-being through better diagnosis and less errors in health care provision that result in readmissions. The implications for Medicare expenditures would depend on the age profile of AI improvements. Alternatively, AI can raise per capita expenditures through increased utilization associated with improved efficacy of care that is insufficient to offset the attendant reduction in price associated with higher productivity.

AI impacts on health-related entitlement expenditures have important dynamic effects through its impact on morbidity and survival rates that evolve over time. If, for example, AI raises survival rates at relatively younger ages but lowers morbidity at older ages, then AI will initially raise expenditures as the survival effect dominates but this increase will be offset by savings from lower utilization at older ages.

To project the effects of AI on entitlement expenditures, we will need current data and projection of the distribution of beneficiaries by age, distribution of Social Security and health care expenditures by age, survival probabilities, and the price of health care services along with elasticities of survival rates, health care expenditures, and health care prices to AI.

Health Care Industry

As before, generative AI is captured as capital-biased technological change. Depending on its complementarities or substitutability with other types of labor in health care, labor demand may rise or fall. Further, AI may well improve the labor productivity of existing workers lowering the cost of the price of health care services, which is set directly or indirectly by the government to recover costs and allow for some profit.

The health care industry hires labor from different occupations *i* with labor compensation in the health care sector denoted *m* at time *t* given by $wl^i_{m,t}$. The health care industry faces a cost of capital r_t^m and rents a capital stock $K_{m,t}$. $Y_{m,t}$ is the quantity of health care services provided that is a function of capital and different labor inputs from industries 1 through J (second equation below). The first equation states that total revenues must equal total costs inclusive of a markup μ.

$$p_{m,t}Y_{m,t} = (1+\mu)C_{m,t}$$
$$Y_{m,t} = F^m(A_{k,t}K_t^m, L_{1,t}^m, \dots, L_{J,t}^m)$$
$$C_{m,t} = r_t^m K_{m,t} + \sum_i wl_{m,t}^i$$

This formulation implies constant passthrough from cost reductions due to AI to health care prices, thereby reducing Medicare expenditures. To determine the effect of AI on the price of health care services, we must make assumptions about the impacts of AI on factor demands and factor costs (i.e., whether wages rise or fall). Imposing functional form assumptions would allow for cost minimizing labor and capital demand to be derived in response to an exogenous increase in capital-biased productivity. In our simulations, we do not simulate changes in capital or wage expenses in the health care industry and instead make direct assumption on the path of health care prices $p_{m,t}$.

Economy-wide Impacts

The government's unified budget constraint is given by the following expression, with outlays on the left-hand side and funding on the right-hand side:

$$G_t + int_t + SS_t + xM_t = \tau^k a_t Y_t + \tau^l (1-a_t) Y_t + def_t$$
$$\Rightarrow G_t + r_t B_t^g + SS_t + xM_t = \tau^k a_t Y_t + \tau^l (1-a_t) Y_t + (B_{t+1}^g - B_t^g)$$

The flow budget constraint states that expenditures on other government spending G_t, interest payments on the debt int_t and entitlements (here Social Security and Medicare) must be financed by taxes levied on capital income, taxes on labor income, and increases in the budget deficit (τ are the respective the tax rates and B_t^g is government debt outstanding). The variable α_t is the (possibly time-varying) capital share of income. The government's flow budget constraint implies a law of motion for the stock of government debt B_t^g that varies with revenue, expenditures and interest costs. This budget constraint consolidates entitlement revenues and expenditures with the overall federal government budget constraint.

Total GDP consists of the sum of value added across all industries i, including health care. Value added is the sum of payments to labor and payments to capital. The relative price of all other industry output is given by $p_{i,t}$ and is assumed to be invariant to the AI shock. The AI shock affects output and revenue via changes in factor income, where $rk_{i,t}$ is capital income for industry i at time t and $wl_{i,t}$ is labor income for industry i at time t:

$$Y_t = \sum_i p_{i,t} Y_{i,t}$$
$$\alpha_t Y_t = \sum_i rk_{i,t}$$

$$(1 - \alpha_t)Y_t = \sum_i wl_{i,t}$$

To estimate the economy-wide impacts of AI on GDP and revenues requires estimates of the impact of AI by industry on the level of capital and labor income and how the distribution between capital and labor income may be impacted within industries. Since AI represents capital-biased technological progress, capital income may increase relative to labor income through higher capital demand and utilization and through displacement of labor for certain occupations.

In our simulations, we do not make assumptions about industry-specific effects of AI on either the capital or labor share. We instead assume effects of AI directly on capital and labor income tax revenue: $\tau^k a_t Y_t$ and $\tau^l (1 - a_t) Y_t$. We also do not model effects of AI on the interest rate but do allow for gradual changes in the interest rate consistent with projections by government forecasters.

4. Methodology

This study aims to present a range of plausible impacts of AI on the federal fiscal outlook. Importantly, these simulations assume an extension in the capability and adoption of AI technology, as opposed to scenarios where AI technology is "frozen" at current levels. For tractability, we depart slightly from the theoretical framework above to simulate the various forms of an AI-driven technological shock on four variables: mortality rates, health care prices, health care utilization, and productivity. The simulation framework captures the connection between AI and federal budgets through four key channels:

1) Mortality rates and longevity: The potentially profound impacts of AI on the efficacy of health care delivery could lead to a sharp decline in age-specific mortality rates and an expansion in longevity. Relative to the previous section, we are considering changes in the path $q_{j,t}$, with mortality effects uniform across age.

2) Health care demand and utilization: Lowering illness burdens and health care needs would lessen the demand for health care utilization in the Medicare population. Alternatively, improved efficacy of health care delivery could raise demand, although such an increase would likely be accompanied by lower prices. Relative to the previous section, we are considering changes in the path $H_{j,t}$, with utilization effects that are uniform across age.

3) Health care prices: As AI lowers overall costs for health care providers and prescription drugs, these savings translate to lower prices for health care services. In addition, the reduction in health care utilization across all age groups, holding supply of health care services constant, could lower health costs in the Medicare program. Relative to the previous section, we are considering changes in the path $p_{m,t}$.

4) Higher aggregate productivity: The above three channels are all unique to health. It may also be the case that the AI revolution also encompasses a more traditional productivity shock, such as the one experienced with widespread adoption of the internet in the late

1990s. A shock of this magnitude would expand tax bases and subsequently raise tax revenue.

The overall fiscal impact of course depends on the relative magnitude of these four channels over time. While certain channels like increased longevity would raise expenditures (by expanding the population of Social Security- and Medicare-eligible beneficiaries), the morbidity and price reductions could help offset these pressures in future periods as the technologies mature. Similarly, higher revenues associated with raised aggregate productivity would offset higher entitlement expenditures.

We begin by modeling changes in longevity by increasing the rate of age- and gender-specific mortality rates as reported by the Social Security Administration (SSA) in the 2023 Old-Age, Survivors, and Disability Insurance Trustees report. The intermediate scenario in the Trustees' projections assumed an annual reduction of 0.74 percent in each age- and gender-specific death rate.[65] In our projections, we increase the annual reduction in mortality rates to either 2 percent or 3 percent and project the subsequent change in the male and female population—beginning with base values in 2023. We do not assume changes in fertility or immigration. This change in longevity serves as the basis for changes in Social Security and Medicare expenditures, the latter of which we apply changes in per capita Medicare spending.

The shocks to mortality that we are predicting fall well within the bounds of historical experience. Japan saw an even larger decline in death rates in the years following the end of WWII. From 1948 through 1958, the unweighted average yearly decline in death rates for each age (0-100) was approximately 6 percent.[66] This is multiple times larger than the death rate reduction we assume in our simulations,[67] although our assumed reduction is substantially higher than the 1.24 percent reduction assumed by the Social Security Trustees in their alternative scenario.[68] Thus, our assumed mortality reduction is higher than envisioned in the Social Security realm of plausible outcomes, but markedly lower than instances of historically fast reductions.

We model changes in health care prices and utilization through the change in per capita Medicare expenditures[69] (which also interacts with changes in the older population). To model changes in Medicare prices, we separate the growth in per capita Medicare spending between changes in price and changes in utilization—first explaining the growth in baseline projected per capita Medicare spending as the product of the growth in prices and growth in utilization. This is performed separately for Medicare Parts A, B, and D, with the growth rates in per capita spending and growth in part-specific costs derived from the 2023 Medicare Trustees Report; the

---

[65] The Board of Trustees, Federal Old-Age and Survivors Insurance and Federal Disability Insurance Trust Funds. 2023. "The 2023 Annual Report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Federal Disability Insurance Trust Funds". https://www.ssa.gov/oact/TR/2023/index.html.

[66] Authors' own calculations; National Institute of Population and Social Security Research. 2023. "Japan, Death Rates (1x1)." https://www.ipss.go.jp/p-toukei/jmd/00/index-en.html.

[67] Importantly, the measure ("average unweighted average yearly decline in death rates for each age") can be compared to "average annual death rate decline" since both annually apply a single death-rate reduction percentage to every age.

[68] The Board of Trustees, Federal Old-Age and Survivors Insurance and Federal Disability Insurance Trust Funds. 2023. "The 2023 Annual Report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Federal Disability Insurance Trust Funds". https://www.ssa.gov/oact/TR/2023/index.html.

[69] Per capita Medicare spending through 2033 is derived from the 2023 Medicare Trustees' Report; we then assume that the annual growth rate for the second decade is equal to the average growth rate of the first decade.

growth rate in baseline utilization is, in effect, the residual value.[70] To model changes in the health care delivery and costs, we adjust the growth rate of prices and utilization for Medicare Part A, B, and D, as specified below. Such an adjustment is highly speculative, and as we alluded to in the literature review, the directional impact of AI-based technological change on these two factors is not even known. Thus, the changes in these parameters are intended to illustrate a potential range of plausible shocks to the health sector. All changes are phased-in uniformly over four years.

We simulate changes in aggregate productivity by increasing the rate of income tax and corporate tax revenues collected as a share of GDP. To scope the size of a potential shock, we look to the productivity shock of the 1990s for guidance, in which the revenue growth rate ultimately grew around 1.5 percent of GDP higher than predicted several years prior, equal to an increased growth rate about 10 percent, prior to the sweeping tax cuts passed at the turn of the century.[71] Using CBO projections as a baseline,[72] for each of the corporate and individual income tax bases, we assume a shock equal to approximately 10 percent of the revenue collected as a share of GDP, or 0.15 percent for corporate taxes and 0.33 percent for individual income taxes. As with Medicare simulations, these shocks are phased in uniformly over four years. We assume no change in the payroll tax base, as an increase would be substantially offset by higher Social Security benefits. In all cases, after a shock has been identified, we apply interest rates as assumed by CBO over the next decade to determine the interest rate changes owing to the change in fiscal expenditures. In each instance, we further assume that average rate of interest payments on the stock of debt over the second decade equals 2.0 percent.

We model four separate scenarios to represent the range of plausible outcomes of the AI-driven shock.

- **Scenario 1, modest reduction in mortality only**: AI modestly improves longevity but does not impact other aspects of the economy. Tax bases are unaffected.
- **Scenario 2, modest efficacy gains**: AI modestly improves longevity through slightly increased utilization and efficacy of care, which also slightly increases price growth. Higher productivity expands tax bases.
- **Scenario 3, major improvement in delivery of care, more efficient delivery**: AI dramatically improves care delivery, with substantial reductions in mortality rates and improvements in efficiency of care. Utilization and prices grow at slower rates relative to the baseline. Higher productivity expands tax bases.
- **Scenario 4, major improvement in delivery of care, more people seek care**: As with the prior scenario, AI dramatically improves care delivery, with attendant reductions in mortality. Utilization increases, which offsets price decreases owing to more efficient care. Higher productivity expands tax bases.

The specific changed are outlined in the table below.

---

[70] The Boards of Trustees, Federal Hospital Insurance and Federal Supplementary Medical Insurance Trust Funds. 2023. “2023 Annual Report of the Federal Hospital Insurance and Federal Supplementary Medical Insurance Trust Funds”. https://www.cms.gov/data-research/statistics-trends-and-reports/trustees-report-trust-funds.

[71] Authors’ own estimation based on projected and realized revenues reported in various CBO fiscal outlooks from the 1990s.

[72] Congressional Budget Office. 2024. “The Long-Term Budget Outlook: 2024 to 2054.” https://www.cbo.gov/publication/59711.

Simulation Parameters

	Longevity	Productivity		Health Care Utilization and Pricing	
	Change in Mortality Rates	Change in Corporate Tax Revenue	Change in Income Tax Revenue	Change in Medicare Utilization (All Parts)	Change in Medicare Price (All Parts)
Scenario					
1	-0.2	0	0	0	0
2	-0.2	0.15	0.33	-0.005	-0.005
3	-0.3	0.15	0.33	-0.01	-0.01
4	-0.3	0.15	0.33	0.01	-0.01

Note: Mortality rate changes are percent reductions in projected age- and gender-specific death rates. Tax revenue changes are increases, as a share of GDP, to respective revenue sources. Utilization and pricing changes are percentage point changes in the annual growth rate.

5. Results

Our results are presented below in various charts and an appendix table. To start, our simulation of a potential increase in longevity shows that such a scenario would increase annual deficits slowly over the 20-year timespan (“budget window”), which naturally makes sense given the cumulative impact of higher annual reductions in mortality rates. Under scenario 1, we see a nominal increase in annual budget deficits of approximately $100 billion after a decade, growing to just over $500 billion by 2044. As a share of GDP, budget deficits would be only modestly affected through the early 2030s, with the negative impact relative to the baseline growing from around 0.2 percent of GDP after a decade to 0.9 percent of GDP in 2044.

Turning to scenario 2, which simulates the impact of modest efficiency gains in health care delivery plus higher tax revenue owing to productivity gains, our simulation projects that such a shock would initially lower deficits relative to the baseline by around 0.8 percent of GDP in the first decade due mainly to revenue and efficiency gains outweighing the impact of mortality reductions. Under this scenario, we assume that reductions in death rates is more modest than in the latter two scenarios, so the impact of an expanding beneficiary population is more muted. As efficiency gains accumulate, annual budget deficits fall by around 1 percent of GDP towards the end of the budget window—roughly $600 billion in nominal terms.

Scenario 3 measures the combined impact of a substantial reductions in mortality rates and a productivity shock which slows the growth of utilization and health care inflation . Here, reductions in death rates are more rapid than in the prior two scenarios. Despite a growing Social Security and Medicare population, revenue gains coupled with reduced per capita Medicare spending drive down deficits relative to the baseline steadily over the budget window. After a decade, deficit reduction totals around 1 percent of GDP. By the end of the budget window, deficit reduction totals around 1.5 percent of GDP—or about $900 billion in nominal terms.

Lastly, scenario 4 investigates the impact of substantial gains in health care delivery efficacy—which in this case drives down health care inflation but leads to greater utilization of health services. Here, with mortality rates falling sharply (as with scenario 3), Social Security expenditures rise markedly throughout the budget window. Medicare spending initially falls as price reductions dominate, but over time the impact of an expanding population and higher utilization more than offset the impact of slower price growth. Year 2037 is the inflection point at which deficit contraction transforms to deficit expansion. By the end of the budget window, annual deficits are around 1 percent of GDP larger relative to the baseline.

Scenario 1 Nominal Impact on Deficit ($b)

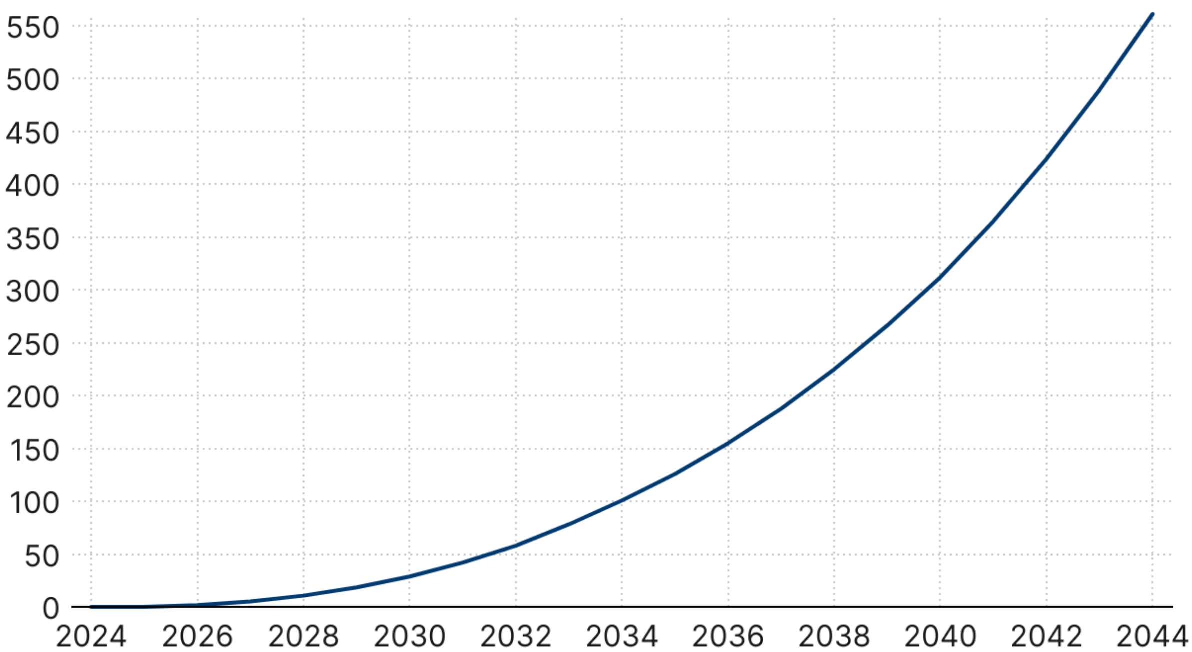

Source: Authors' calculations

BROOKINGS

Scenario 1 vs Baseline Deficit (% of GDP)

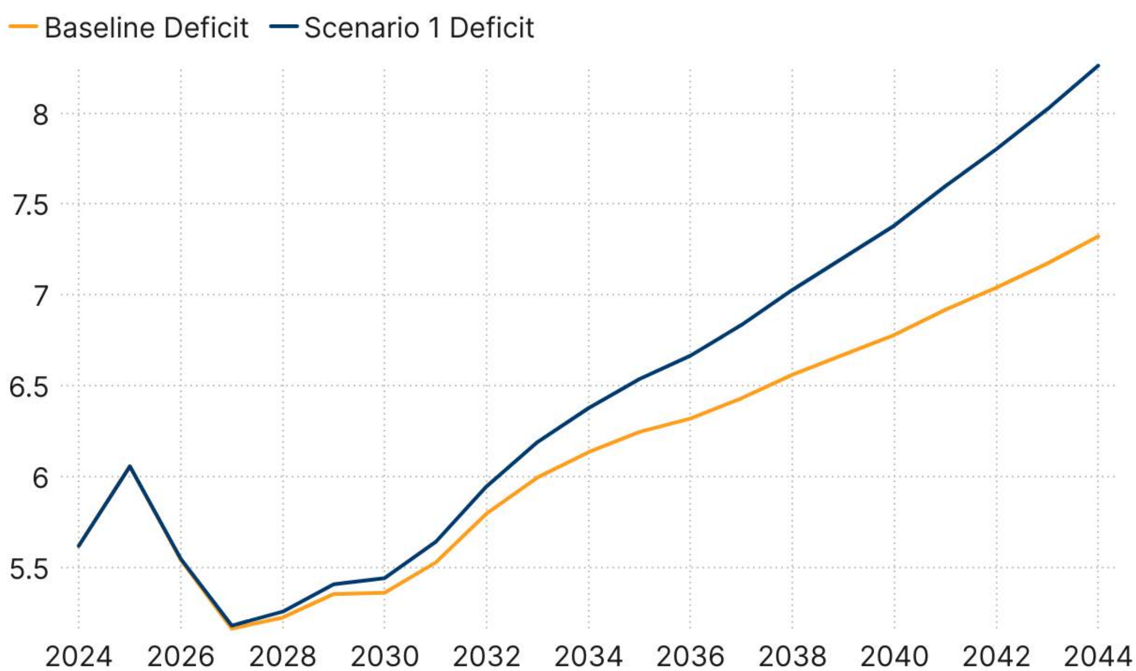

Source: Authors' calculations

Center on Regulation and Markets at BROOKINGS

Scenario 2 Nominal Impact on Deficit ($b)

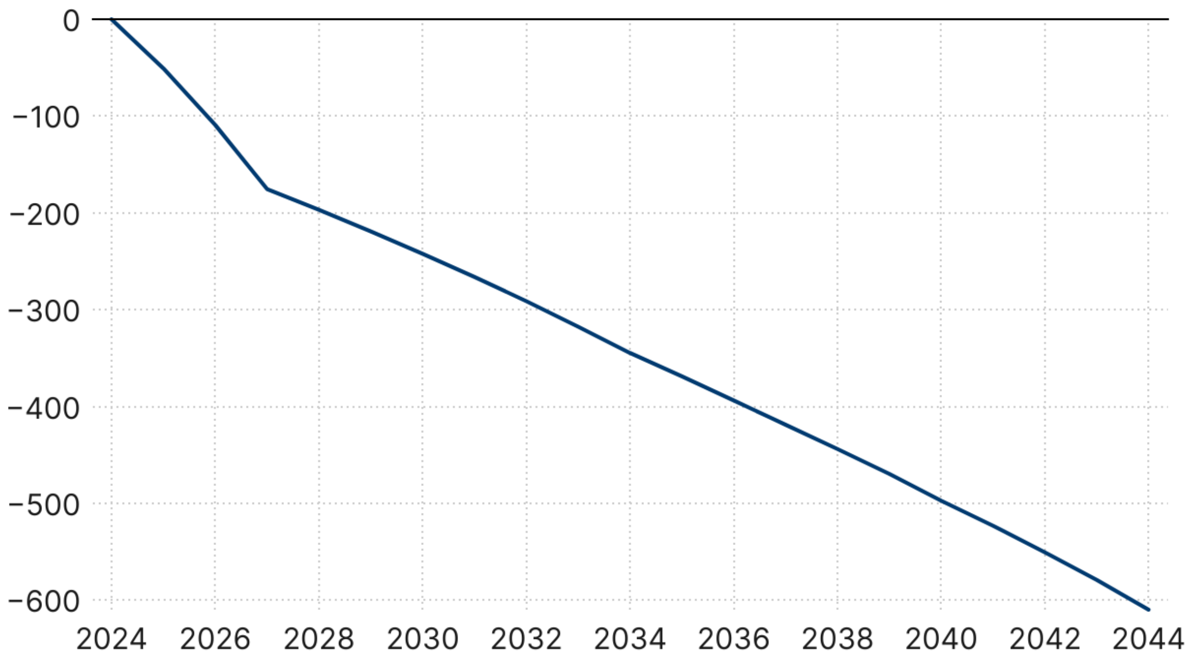

Source: Authors' calculations

BROOKINGS

Scenario 2 vs Baseline Deficit (% of GDP)

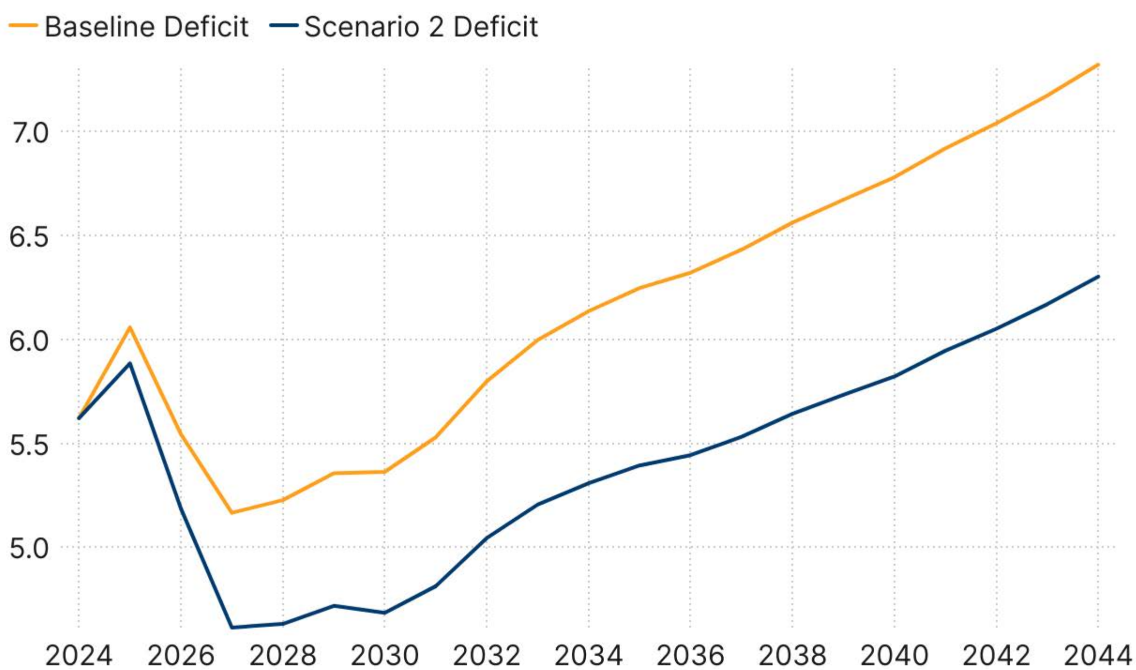

Source: Authors' calculations

B | Center on Regulation and Markets at BROOKINGS

Scenario 3 Nominal Impact on Deficit ($b)

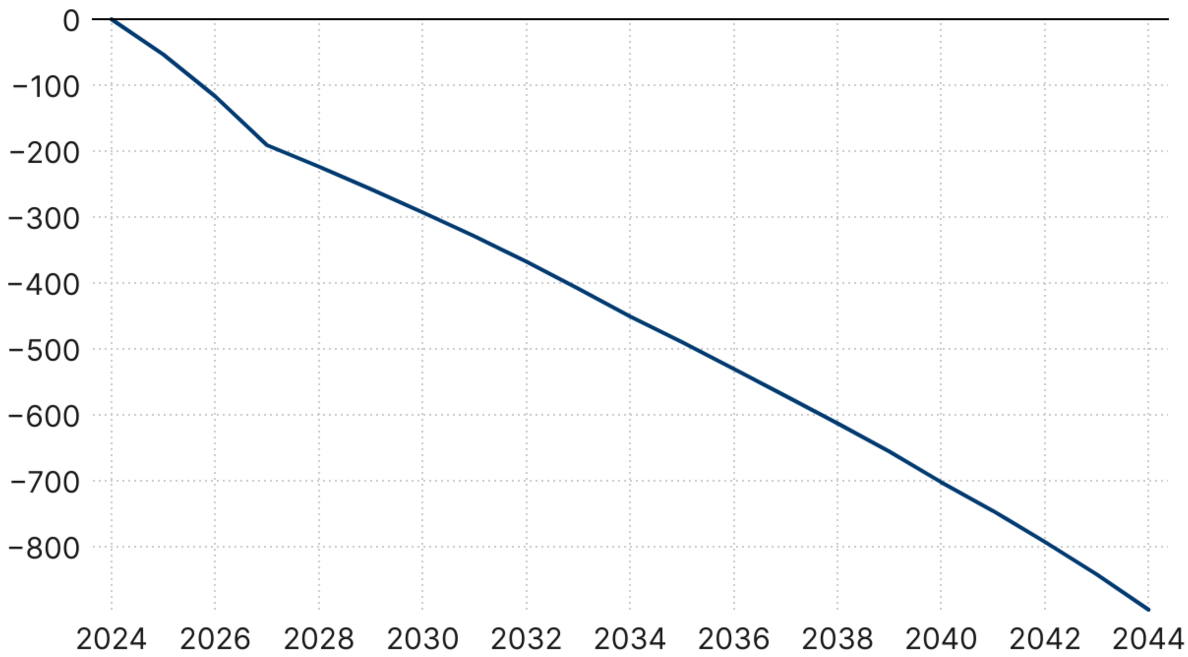

Source: Authors' calculations

BROOKINGS

Scenario 3 vs Baseline Deficit (% of GDP)

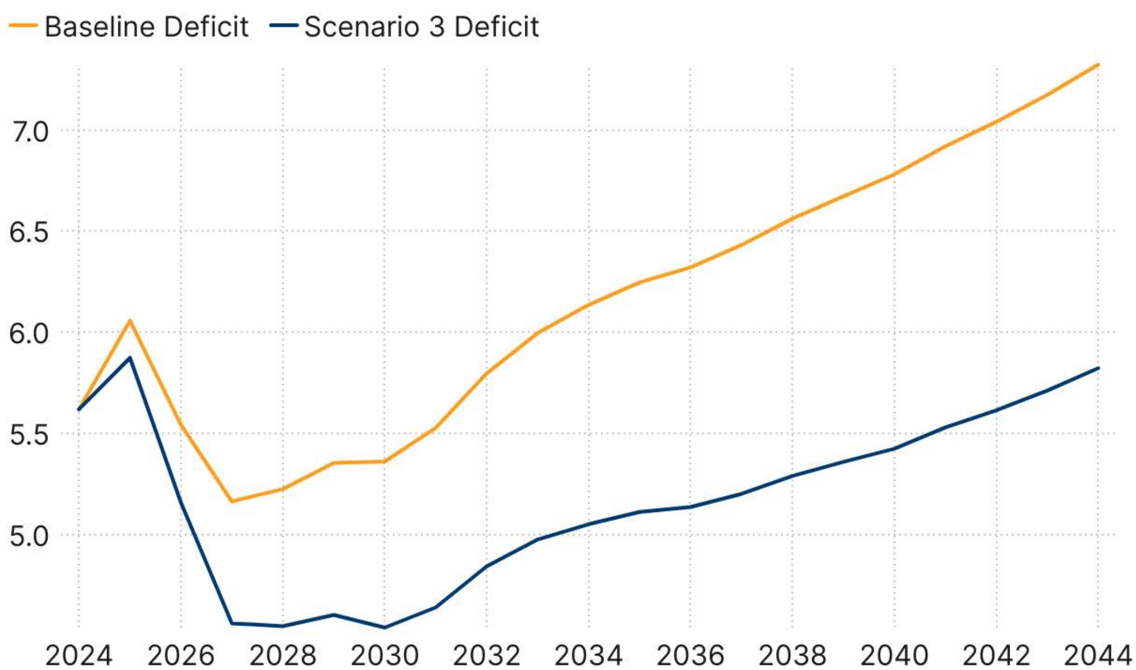

Source: Authors' calculations

B Center on Regulation and Markets at BROOKINGS

Scenario 4 Nominal Impact on Deficit ($b)

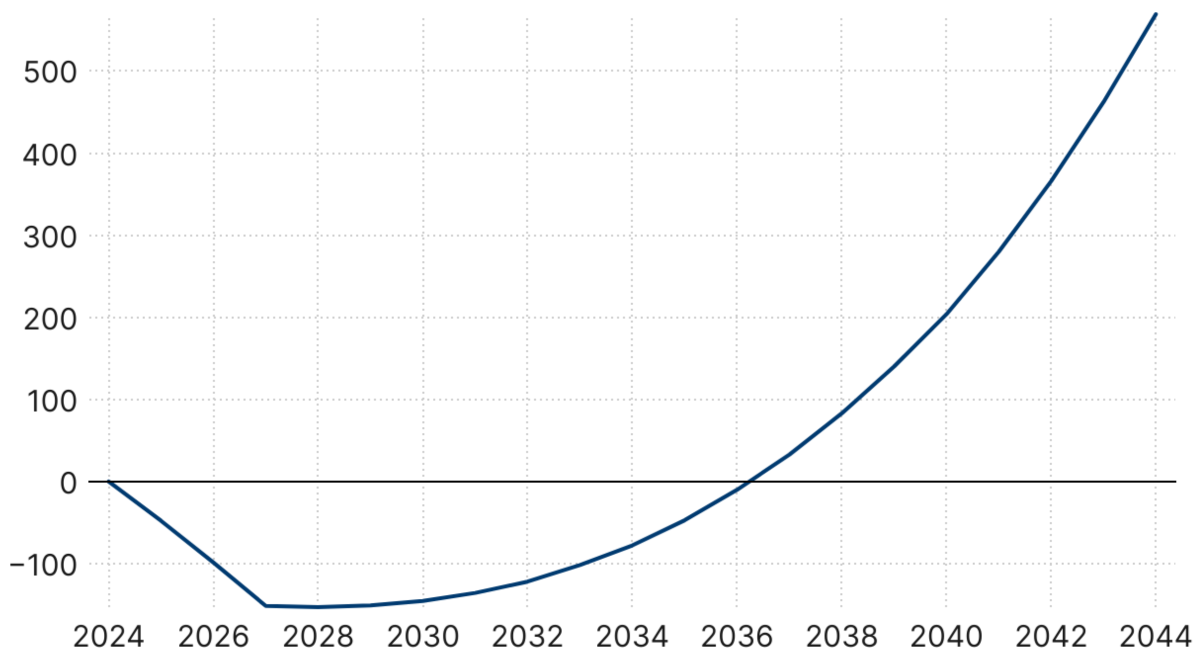

Source: Authors' calculations

BROOKINGS

Scenario 4 vs Baseline Deficit (% of GDP)

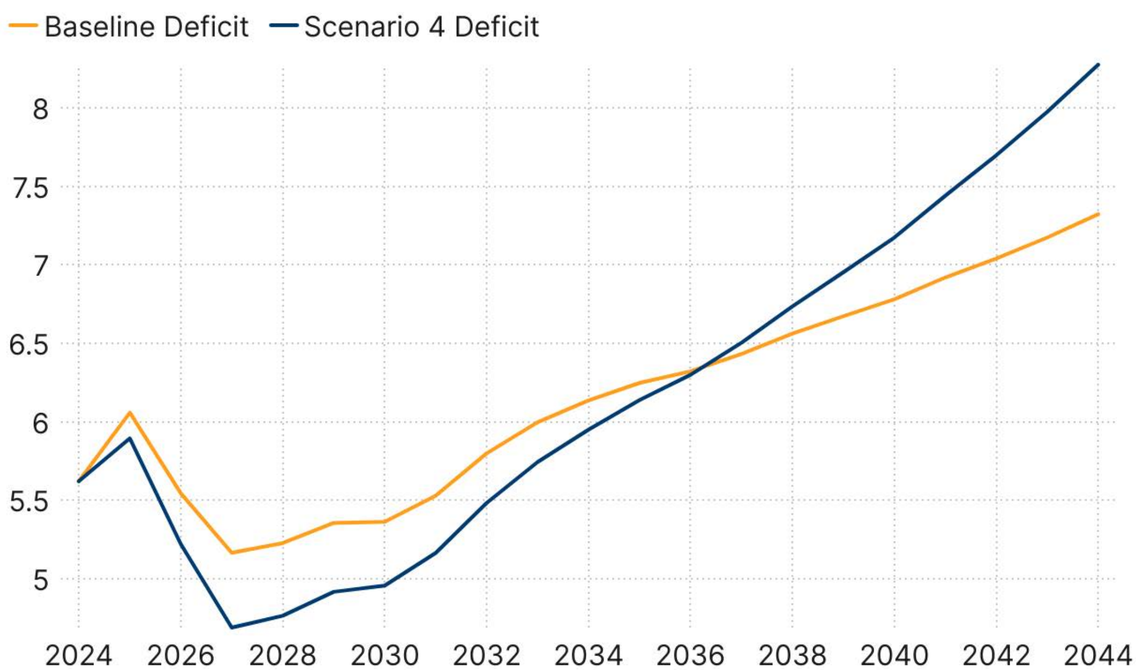

Source: Authors' calculations

B Center on Regulation and Markets at BROOKINGS

6. Conclusion

AI offers a massive opportunity to transform the economy, perhaps especially in the realm of health care delivery. To date, much of the attention paid to AI has been focused on labor productivity, social implications, and corporate profits. We believe the lack of attention paid to health care pricing, efficacy, and utilization is an oversight; AI has already demonstrated substantial progress and enormous potential in improving health care outcomes.

We simulate four separate shocks to the federal budget, each representing a different type of shock. The consensus among economists, at least informally, appears to be that AI will improve the budget outlook due to higher revenues. This confidence may be misplaced. A shock to revenues that is in-line with the experience of the 1990s would meaningfully improve the fiscal outlook but would be offset by other health-related factors. The real promise in addressing our nation's longstanding fiscal imbalance lies in AI's ability to bend the cost curve for growth in per capita Medicare spending.

In this paper, we perform representative simulations to show that the impact of AI on the fiscal outlook depends critically on changes to a handful of factors: mortality, health care inflation, health care utilization, and productivity. Depending on how AI affects these various factors, the plausible impacts on annual deficits range from an increase of just under 1 percent of GDP to a decrease of just over 1.5 percent of GDP—exceeding in absolute value terms the impact of the productivity shock of the late 1990s.

7. Appendix

Scenario Impacts on Deficit, Nominal and %GDP

	Impact on deficit ($b)				Impact on deficit (%GDP)			
Year	Scenario 1	Scenario 2	Scenario 3	Scenario 4	Scenario 1	Scenario 2	Scenario 3	Scenario 4
2024	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2025	0.00	−50.74	−53.58	−47.84	0.00	0.17	0.18	0.16
2026	1.68	−109.12	−116.96	−98.56	−0.01	0.36	0.38	0.32
2027	5.21	−175.39	−191.08	−151.44	−0.02	0.55	0.60	0.48
2028	10.64	−196.74	−223.52	−152.88	−0.03	0.60	0.68	0.46
2029	18.45	−219.12	−257.82	−150.91	−0.05	0.64	0.75	0.44
2030	28.78	−242.28	−292.99	−145.38	−0.08	0.68	0.82	0.41
2031	41.96	−266.09	−329.22	−135.61	−0.11	0.72	0.89	0.36
2032	57.87	−291.16	−367.65	−121.84	−0.15	0.75	0.95	0.32
2033	78.13	−317.54	−408.57	−101.51	−0.19	0.79	1.02	0.25
2034	100.76	−344.69	−450.75	−77.67	−0.24	0.83	1.08	0.19
2035	125.82	−368.82	−489.79	−47.09	−0.29	0.85	1.13	0.11
2036	154.89	−393.69	−530.44	−10.05	−0.35	0.88	1.18	0.02
2037	187.71	−418.78	−571.59	32.98	−0.40	0.90	1.23	−0.07
2038	224.92	−443.92	−613.02	83.09	−0.47	0.92	1.27	−0.17
2039	266.44	−469.59	−655.64	140.15	−0.53	0.94	1.31	−0.28
2040	312.02	−497.28	−702.23	203.72	−0.60	0.96	1.35	−0.39
2041	365.21	−523.14	−745.91	280.05	−0.68	0.97	1.39	−0.52
2042	424.22	−550.60	−792.91	365.96	−0.76	0.99	1.42	−0.66
2043	489.47	−579.12	−842.20	462.20	−0.85	1.00	1.46	−0.80
2044	561.24	−609.77	−895.75	569.19	−0.94	1.02	1.50	−0.95

Source: Authors' Calculations

Center on Regulation and Markets at BROOKINGS

The Center on Regulation and Markets at Brookings provides independent, non-partisan research on regulatory policy, applied broadly across microeconomic fields. It creates and promotes independent economic scholarship to inform regulatory policymaking, the regulatory process, and the efficient and equitable functioning of economic markets.

Questions about the research? Email communications@brookings.edu.
Be sure to include the title of this paper in your inquiry.