

For release on delivery  
8:50 a.m. EDT  
September 22, 2023

Generative AI, Productivity, the Labor Market, and Choice Behavior

Remarks by

Lisa D. Cook

Member

Board of Governors of the Federal Reserve System

at the

National Bureau of Economic Research  
Economics of Artificial Intelligence Conference, Fall 2023

Toronto, Canada

September 22, 2023

## Introduction

As I was in 2018, I am excited to speak to you at the National Bureau of Economic Research (NBER) artificial intelligence (AI) conference, in a city that is a world-class center of AI research and business start-ups, to discuss AI and its prospective effects on productivity and the labor market. Outside of those of us who have spent many years researching the economics of innovation, it seems that AI is having a moment. The surge in excitement and trepidation about AI is palpable. Google searches for “AI” have tripled worldwide since 2022, fueled by the buzz about ChatGPT. Of course, this group saw it coming as early as 2017, when the first NBER AI conference was held here in Toronto, and many of you saw it coming much earlier than that.

I will focus my remarks on generative AI, which creates new content largely in response to natural language prompts.<sup>1</sup> As this audience knows, image and text classification—discriminative AI—has been in use for many years and is remarkably effective. I have used it to identify demographic characteristics of entrepreneurs in my own research.<sup>2</sup> In contrast, effective generative AI is a very recent development and seems to be a leap forward into something new. Applications of generative AI range from the prosaic, like reducing the monotony of writing routine memos, to the wondrous, like protein structure prediction and drug discovery.

---

<sup>1</sup> The views I express here are my own and not those of the Board of Governors of the Federal Reserve System or the Federal Open Market Committee.

<sup>2</sup> See Lisa D. Cook, Matt Marx, and Emmanuel Yimfor (2022), “Funding Black High-Growth Startups,” NBER Working Paper Series 30682 (Cambridge, Mass.: National Bureau of Economic Research, November), <https://www.nber.org/papers/w30682>.

Of course, experts emphasize that at their core, all forms of AI are an exercise in prediction, and technically that is true.<sup>3</sup> To the layperson, though, a chatbot that is nearly good enough to pass the Turing test is substantially different from the U.S. Postal Service using AI to read your handwriting. Some of the uses of generative AI may be unsettling. For example, concerns about the ability of generative AI to impersonate individuals to harm their reputation or violate their privacy exist and are growing. Moreover, observers have noted that AI models sometimes harbor, if not amplify, the biases found in their training data, leading to malign effects on decisions about mortgage approvals, insurance rates, medical diagnoses, and even pretrial detention.<sup>4</sup> And discrimination is not just an equity issue—it also holds down economic growth, as I show in my own work.<sup>5</sup>

The range of potential social effects of AI is wide, as will be explored in the next presentation.<sup>6</sup> In general, I am optimistic about broad benefits accruing to the economy and society from the use of generative AI—including more productive and less tedious work in offices, labs, factories, and warehouses—provided we address the very real concerns I just mentioned, and others like them.

As we consider how to foster the emerging benefits of AI and guard against unwelcome harms, it is important to keep in mind that the path from innovation to greater

---

<sup>3</sup> On AI as prediction, see Ajay Agrawal, Joshua S. Gans, and Avi Goldfarb (2019), “Artificial Intelligence: The Ambiguous Labor Market Impact of Automating Prediction,” *Journal of Economic Perspectives*, vol. 33 (Spring), pp. 31–50.

<sup>4</sup> On criminal justice, see Laurel Eckhouse, Kristian Lum, Cynthia Conti-Cook, and Julie Ciccolini (2019), “Layers of Bias: A Unified Approach for Understanding Problems with Risk Assessment,” *Criminal Justice and Behavior*, vol. 46 (February), pp. 185–209. For a broader discussion of bias mitigation in AI, see Xavier Ferrer, Tom van Nuenen, Jose M. Such, Mark Coté, and Natalia Criado (2021), “Bias and Discrimination in AI: A Cross-Disciplinary Perspective,” *IEEE Technology and Society Magazine*, vol. 40 (June), pp. 72–80.

<sup>5</sup> See Lisa D. Cook, Janet Gerson, and Jennifer Kuan (2022), “Closing the Innovation Gap in Pink and Black,” *Entrepreneurship and Innovation Policy and the Economy*, vol. 1, pp. 43–66.

<sup>6</sup> See Charles I. Jones (2023), “The A.I. Dilemma: Growth versus Existential Risk,” working paper, September 12.

welfare passes through the choices of individuals in a social context—in the corner office, in government, and in the minds of workers and consumers—and progress could stall or accelerate in any of these places. I will return to this point later after offering some thoughts on the potential for AI to affect productivity and the labor market.

Why do I focus on AI as a monetary policymaker? The Federal Reserve’s dual mandate is to promote maximum employment and stable prices. When firms deploy technologies that make workers more productive, they create the conditions for greater wage growth consistent with stable prices. And the labor market adjustment that follows as the economy adapts to technical change can affect maximum employment.

### **AI and Productivity**

The impact of AI on the economy and monetary policy will depend on whether AI is just another app or something more profound. The most consequential innovations in the past have been general purpose technologies that have broadly transformed the economy over an extended period of time. We are living through the ongoing transformation fueled by electronic information technology, for example, and electrification had a similar effect in the early 20th century. General purpose technologies have three key features: (1) they are widely used across the economy, (2) they improve steadily over a long period of time, and (3) they raise the productivity of research and development (R&D).<sup>7</sup> Could generative AI have these features? I will consider each in turn.

---

<sup>7</sup> For the definition of general purpose technology, see Boyan Jovanovic and Peter L. Rousseau (2005), “General Purpose Technologies,” in Philippe Aghion and Steven N. Durlauf, eds., *Handbook of Economic Growth*, vol. 1B (Amsterdam: Elsevier), pp. 1181–224.

First, is generative AI widely used? It is easy to see the potential, and we seem to be headed for widespread use. Generative AI makes communication more efficient, and nearly all human activities—and all industries—involve communication. It is true that if you let generative AI draft an email, write the minutes of a meeting, or research a topic, you will have to review, fact-check, and edit the result. Nonetheless, thanks to AI's contribution, you may be much closer to your goal when you start than if you began with a blank page. Empirical evidence is still patchy, but there is work showing that generative AI improves productivity in a variety of settings, including computer coding, customer service, language translation, and robotics.<sup>8</sup>

Second, will AI itself improve steadily over time? If we look backward, we can see that although the history of the computer language models at the core of generative AI goes back at least to the 1950s, there has been an explosion of technical progress in very recent years as LLMs, or large language models, using neural networks have emerged. Whether that explosive progress can be sustained is an open question, although the concerted efforts here in Toronto and elsewhere bode well for continued innovation. To draw an analogy, the sustained progress in solid-state electronics correctly predicted by Gordon Moore in 1965 looks like a law from a distance. But, in reality, each new

---

<sup>8</sup> Empirical papers providing evidence of the effectiveness of generative AI span a variety of tasks, including financial services, computer programming, call centers, and writing. For financial services, see Tracy Yang, Tian Lu, Beibei Li, and Lu Xianghua (2020), "Personalizing Debt Collections: Combining Reinforcement Learning and Field Experiment," ICIS 2020 Proceedings, December 14. For computer programming, see Sida Peng, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer (2023), "The Impact of AI on Developer Productivity: Evidence from GitHub Copilot," unpublished paper, Cornell University, arXiv, February, <https://arxiv.org/pdf/2302.06590.pdf>. For call centers, see Erik Brynjolfsson, Danielle Li, and Lindsay R. Raymond (2023), "Generative AI at Work," NBER Working Paper Series 31161 (Cambridge, Mass.: National Bureau of Economic Research, April), <https://www.nber.org/papers/w31161>. For writing, see Shakked Noy and Whitney Zhang (2023), "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence," working paper, March 10; Emma van Inwegen, Zanele T. Munyikwa, and John J. Horton (2023), "Algorithmic Writing Assistance on Jobseekers' Resumes Increases Hires," NBER Working Paper Series 30886 (Cambridge, Mass.: National Bureau of Economic Research, January), <https://www.nber.org/papers/w30886>.

generation of chip technology represents the coordinated effort of hundreds of scientists and engineers solving seemingly intractable problems.<sup>9</sup> Continuing advances in model architecture, data curation, and computation will be essential for the continual improvement of AI models and implementation.

Third, does generative AI make R&D more productive? Some potential for efficiency improvements in the scientific process when it comes to literature review and writing is obvious. Yet AI can go much deeper, discovering patterns in data and in previous research to generate hypotheses for testing that may not have occurred to researchers. Work by Ludwig and Mullainathan on exactly this topic will be presented shortly.

All told, generative AI seems promising as a general purpose technology. Of course, you will get a much deeper dive into this question later this morning with the Eloundou, Manning, Mishkin, and Rock presentation. In their work, they find that 80 percent of the U.S. workforce will see at least some of their tasks transformed by generative AI. The authors of that paper do not take a stand on how fast this transformation will take place. Nor will I. However, we do know that historically the journey from innovation to productivity has sometimes been a long and uneven one. An often-cited example is the electric dynamo, which was first used in the U.S. in the 1890s but did not boost manufacturing productivity until the 1920s.<sup>10</sup> Things now are a bit

---

<sup>9</sup> See Hassan N. Khan, David A. Hounshell, and Erica R.H. Fuchs (2018), “Science and Research Policy at the End of Moore’s Law,” *Nature Electronics*, vol. 1 (January), pp. 14–21. Moore’s initial prediction was later revised to state that the number of components on a cost-effective integrated circuit would double every two years. See Gordon E. Moore (1965), “Cramming More Components onto Integrated Circuits,” *Electronics*, vol. 38(8), pp. 114–17; Gordon E. Moore (1975), “Progress in Digital Integrated Electronics,” *Technical Digest 1975, International Electron Devices Meeting, IEEE*, pp. 11–13.

<sup>10</sup> See Paul A. David (1990), “The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox,” *American Economic Review*, vol. 80 (May), pp. 355–61.

more hopeful than that example suggests, though: The lag between invention and adoption has been substantially reduced since the 19th century.<sup>11</sup> Adoption of generative AI is certainly happening at a rapid clip. Even so, the full benefit of a technology only follows adoption when suitable complementary investments have been made.<sup>12</sup> These can include changes in corporate structure and management practices, worker training, and the adjustment of the mix of capital in use. On the last point, we may have a head start, as AI will be deployed in a world with a massive stock of information technology already in place. New business formation will surely play a role as well, as historically much of productivity growth has followed from the entry of firms starting with a clean slate—and the exit of firms that were slow to adapt.<sup>13</sup>

### **Labor Market Effects**

As with all revolutionary technologies, when we turn our attention from productivity to the labor market, many express concern, focusing on jobs that may disappear, while others focus on which jobs will replace them. Economic history suggests cautious optimism here. When the world switched from horse-drawn transport to motor vehicles, jobs for stable hands disappeared, but jobs for auto mechanics took their place.<sup>14</sup> New technologies may displace some types of labor, but they can also raise the productivity and incomes of jobs they create or complement. The increase in

---

<sup>11</sup> See Diego Comin, Danial Lashkari, and Martí Mestieri (2021), “Structural Change with Long-Run Income and Price Effects,” *Econometrica*, vol. 89 (January), pp. 311–74.

<sup>12</sup> See Erik Brynjolfsson, Daniel Rock, and Chad Syverson (2021), “The Productivity J-Curve: How Intangibles Complement General Purpose Technologies,” *American Economic Journal: Macroeconomics*, vol. 13 (January), pp. 333–72.

<sup>13</sup> See Ryan Decker, John Haltiwanger, Ron Jarmin, and Javier Miranda (2014), “The Role of Entrepreneurship in US Job Creation and Economic Dynamism,” *Journal of Economic Perspectives*, vol. 28 (Summer), pp. 3–24.

<sup>14</sup> See Georgios Petropoulos (2018), “The Impact of Artificial Intelligence on Employment,” in Max Neufeind, Jacqueline O’Reilly, and Florian Ranft, eds., *Work in the Digital Age* (London: Rowman & Littlefield International), pp. 119–32.

consumption that follows may raise demand for labor overall. Nonetheless, the displacement effect might be concentrated and the productivity effect more diffuse. Therefore, while many workers throughout the economy benefit, a smaller set bear the brunt of the negative effects. Just as the introduction of computerized machine tools replaced skilled machinists and personal computers made many routine clerical and administrative jobs obsolete, the widespread adoption of AI will be a difficult transition for some workers.<sup>15</sup>

But the labor market effects of technological change are more subtle than just creating and eliminating positions. Labor economists encourage us to think of work in terms of tasks, not jobs.<sup>16</sup> As firms rethink their product lines and how they produce their goods and services in response to technical change, the composition of the tasks that need to be performed changes. Here, the portfolio of skills that workers have to offer is crucial. Can you shift to a new position that requires a different mix of your skills? For workers with a diverse skill set, and for workers with broad skills, like critical thinking and project management, the answer may well be “yes.” For others, like the stable hand who was highly skilled in grooming horses, the answer may be “no.”<sup>17</sup>

---

<sup>15</sup> See Leah Platt Boustan, Jiwon Choi, and David Clingingsmith (2022), “Automation after the Assembly Line: Computerized Machine Tools, Employment and Productivity in the United States,” NBER Working Paper Series 30400 (Cambridge, Mass.: National Bureau of Economic Research, August; revised October), [https://www.nber.org/system/files/working\\_papers/w30400/w30400.pdf](https://www.nber.org/system/files/working_papers/w30400/w30400.pdf); David H. Autor and David Dorn (2013), “The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market,” *American Economic Review*, vol. 103 (August), pp. 1553–97.

<sup>16</sup> See David H. Autor, Frank Levy, and Richard J. Murnane (2003), “The Skill Content of Recent Technological Change: An Empirical Exploration,” *Quarterly Journal of Economics*, vol. 118 (November), pp. 1279–333; Daron Acemoglu and Pascual Restrepo (2018), “The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment,” *American Economic Review*, vol. 108 (June), pp. 1488–542.

<sup>17</sup> On the difficulty faced by workers with eliminated jobs, see David Hummels, Jakob R. Munch, Lars Skipper, and Chong Xiang (2012), “Offshoring, Transition, and Training: Evidence from Danish Matched Worker-Firm Data,” *American Economic Review*, vol. 102 (May), pp. 424–28.



The ability of workers to move to where they are needed as the task composition of production changes will also be an important determinant of how successfully the economy adapts to the new jobs created in response to AI. For example, how quickly will education and training react to the market signals of the skills that are needed? How will AI affect the range of skills required within firms and how will firms restructure in response? And how efficiently will the labor market match job seekers to suitable vacancies?

While the Federal Reserve does not have a role in setting policies to help workers directly, I do not want to suggest that this transition will be easy or painless. Any large change in the labor force will generate disruptions and challenges that will need to be addressed to help workers adapt and thrive. The benefit of AI to society as a whole will depend on the adaptability of workers' skills, how well they are retrained or redeployed, and how policymakers choose to support the groups that are hardest hit by these changes.

### **Choice Behavior**

The potential for far-reaching changes to the economy from generative AI is clear, but the pace and extent of the changes will depend on the choices made by workers, managers, and policymakers. AI makes predictions, but AI does not make choices. Ultimately, human beings are still in control.

For workers, preparing for the AI-enhanced economy is a tricky task. What should students focus on in school? What college and university courses should be developed and mandatory? What kinds of continuing education are needed? It is safe to say that generative AI will make knowledge work more efficient—a worker can do more research, communication, design, and the like in a day. And, while some observers might

warn that means fewer such workers, it is more likely we will need more of them. After all, when knowledge workers can accomplish more in an hour, firms have an incentive to use more of them, not fewer. So the demand for STEM skills will be robust, as it has been throughout the information age, but AI technology may strengthen the rising demand for social skills as well.<sup>18</sup> Some of the job titles will be brand new. A search for “generative AI” jobs on Indeed.com early this week found over 2,000 listings, including such titles as “prompt engineer” and “newsroom generative AI lead.”

Among firms, success deploying AI will depend on strategic decisions, such as investing in training, reorganization, and product development. Financing will need to be available to existing firms that appear to best leverage the potential of AI and to the innovative new firms that will surely appear with AI-based business models.

Policymakers, too, at all levels of government, will have to confront the changing world. Importantly, in the policy arena—as well as health care, consumer finance, insurance, and many others—decisionmakers have legal and ethical duties to be deliberate about the effects their choices have on affected groups. In this context, an AI black box with no insight into the decision-making process is of limited value. As a policymaker, I look upon model-generated forecasts with a skeptical eye, if they are not coupled with a plausible explanation for the driving factors behind them. More generally, when stakeholders have an opportunity to appeal a decision, they are entitled to understand how the decision was made—an issue I emphasized when I spoke at the

---

<sup>18</sup> See David J. Deming (2017), “The Growing Importance of Social Skills in the Labor Market,” *Quarterly Journal of Economics*, vol. 132 (November), pp. 1593–640.

2018 meeting of this group.<sup>19</sup> So I am particularly interested in seeing progress on “explainable AI,” which may help bridge the divide between the technical sphere and the user.<sup>20</sup>

In short, the impact of generative AI, as with all technical change, has to be understood in terms of human choice behavior in specific social and institutional contexts. Generative AI will change the choice set available to consumers, firms, and policymakers. As it happens, because economists study choice behavior, we are well positioned to contribute to the debate about AI and welfare and to anticipate the trajectory of this exciting trend. Some questions you might consider include: Are there ways to limit the labor-force disruptions of AI while capturing its job-creating potential? What new training and skill development will be needed to capture AI’s benefits? Can productivity measures be improved to better capture how quickly AI is affecting the economy?

I am delighted that this group continues to produce valuable, insightful work on these and other questions on the economics of AI, including the many interesting, thoughtful papers to be presented today.

Thank you.

---

<sup>19</sup> See Lisa Cook (2018), “The Consequences of AI-Based Decisions or Outsourcing Tasks When Humans Fail,” presented at the NBER Economics of Artificial Intelligence Conference 2018, Toronto, September 13–14.

<sup>20</sup> See Pantelis Linardatos, Vasilis Papastefanopoulos, and Sotiris Kotsiantis (2021), “Explainable AI: A Review of Machine Learning Interpretability Methods,” *Entropy*, 23, 18.