THE MACROECONOMICS OF

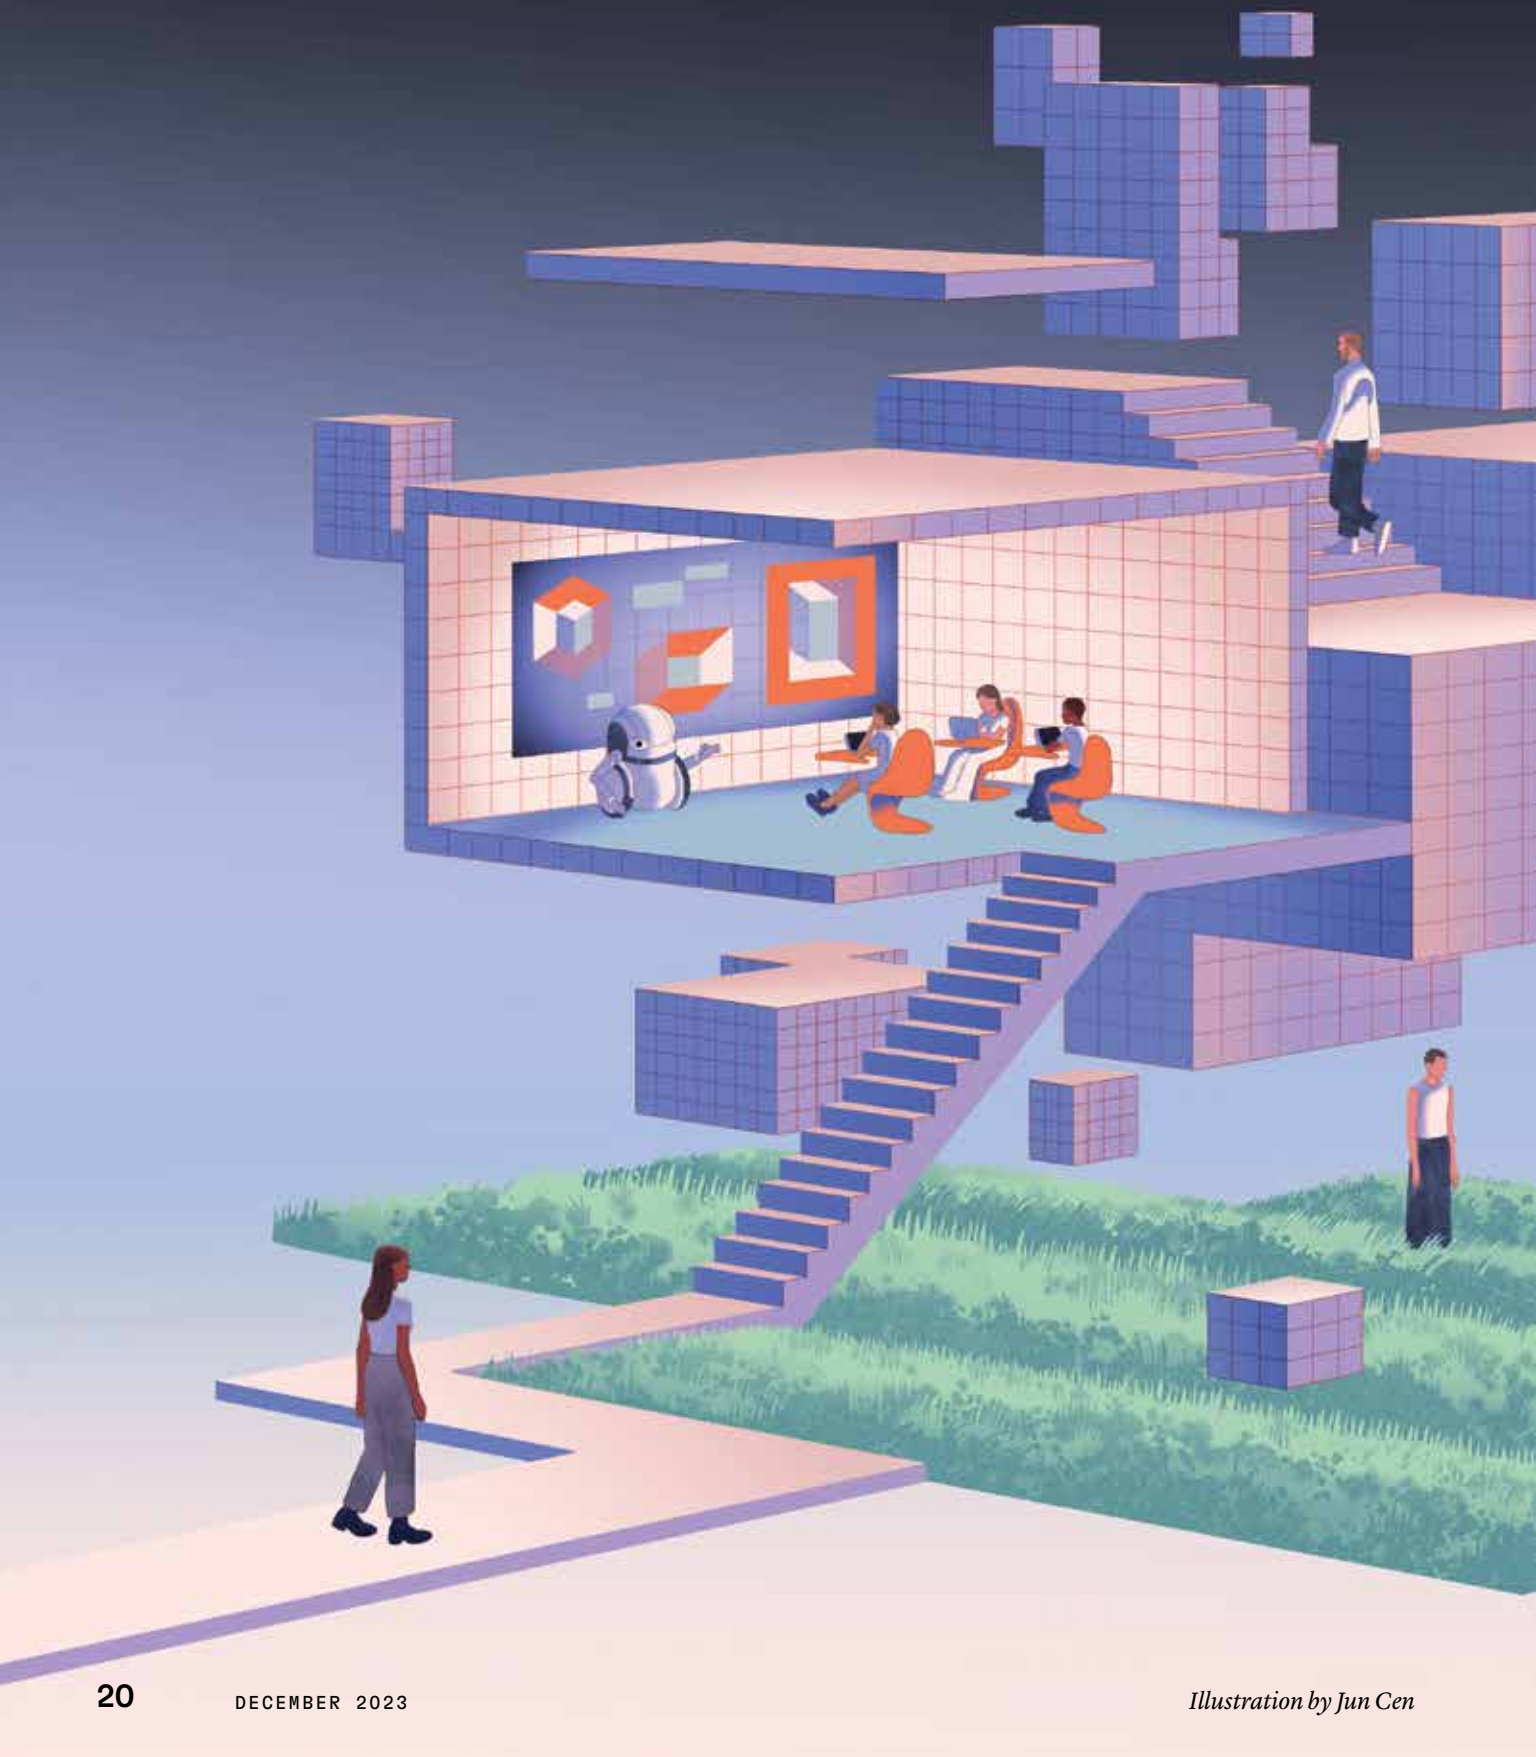

Illustration by Jun Cen

ARTIFICIAL INTELLIGENCE

The collective decisions we make today will determine how AI affects productivity growth, income inequality, and industrial concentration

Erik Brynjolfsson and Gabriel Unger

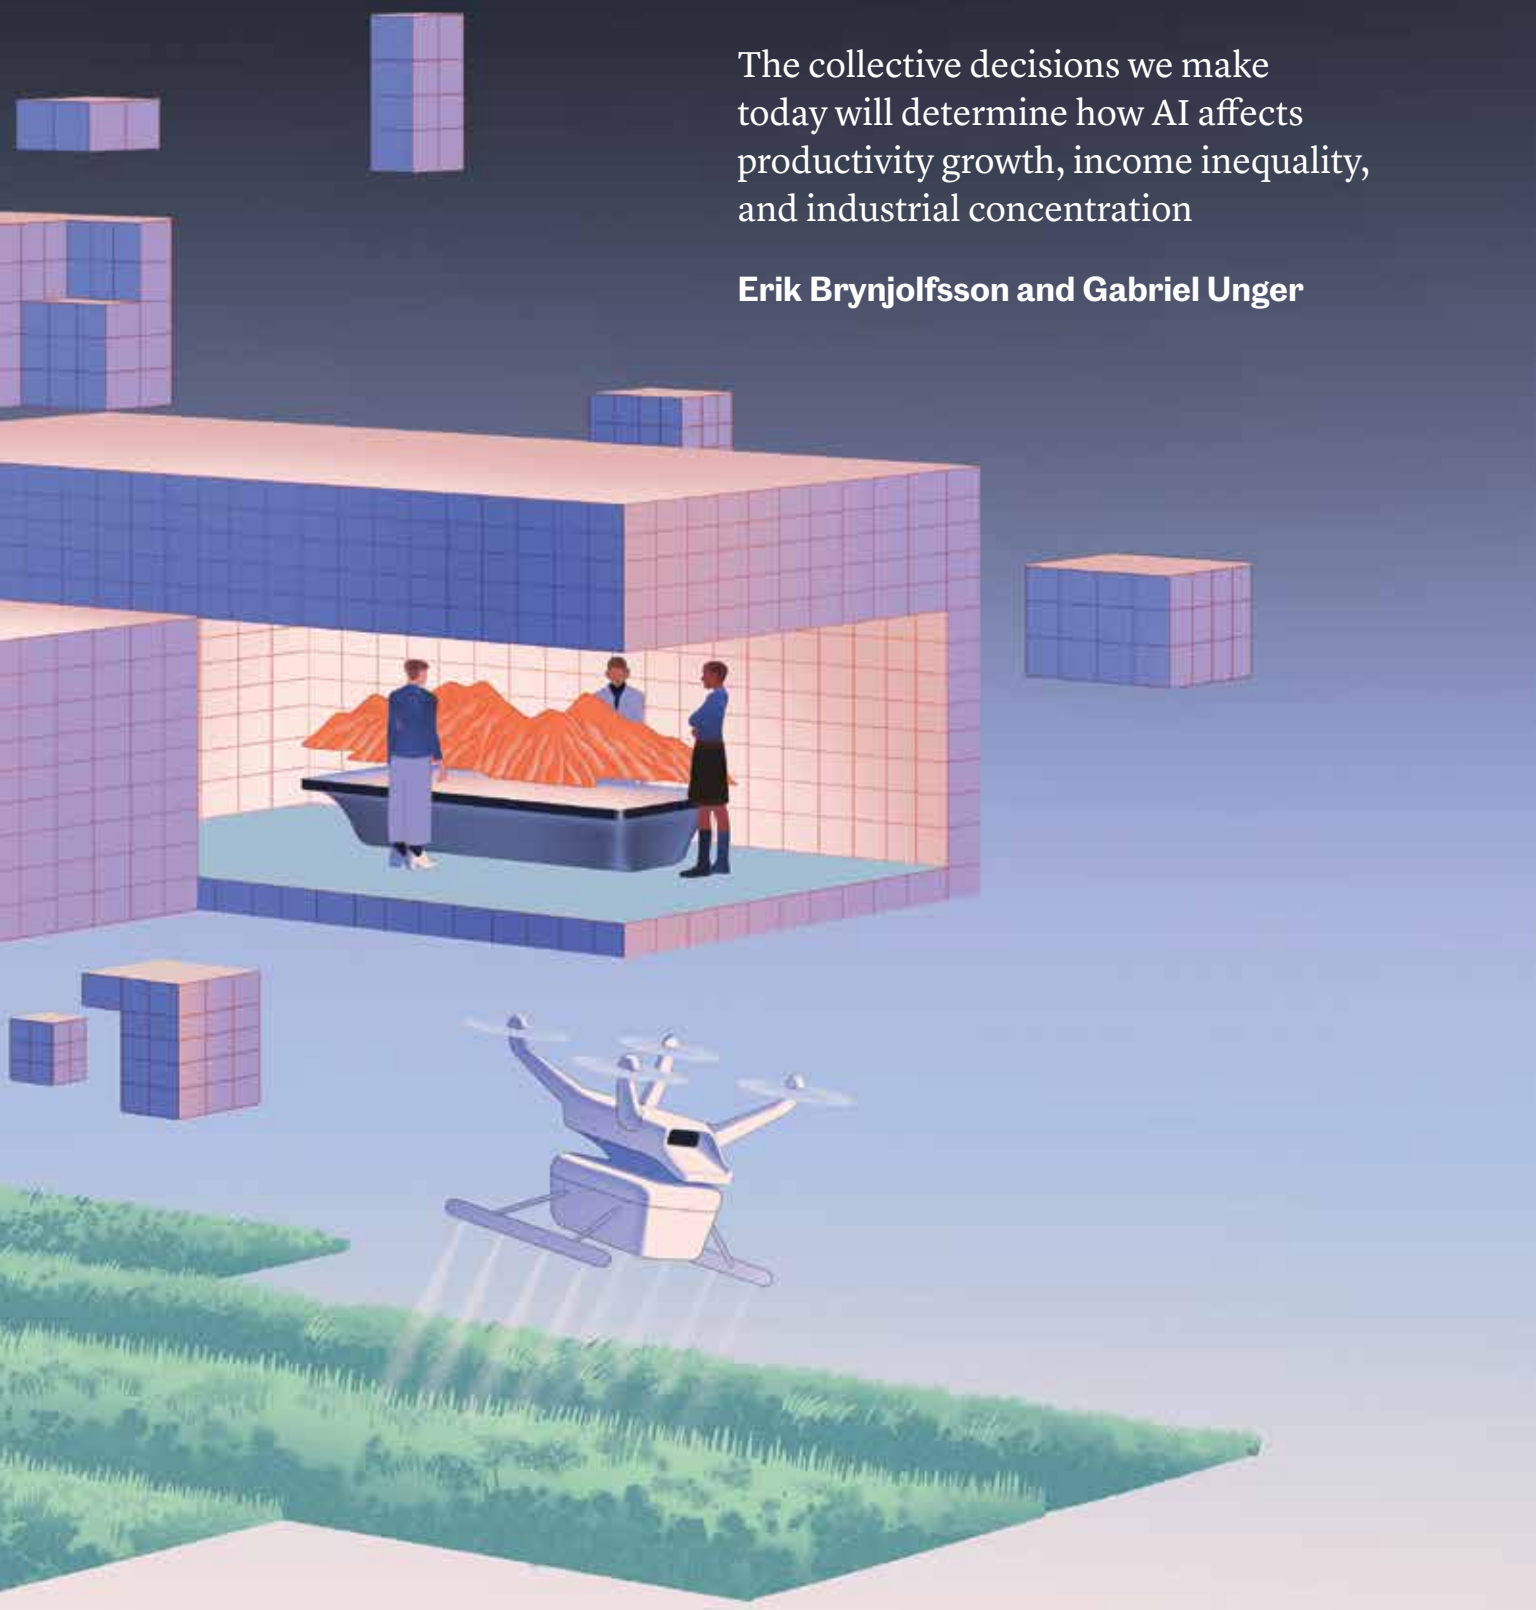

Economists have a poor track record of predicting the future. And Silicon Valley repeatedly cycles through hope and disappointment over the next big technology. So a healthy skepticism toward any pronouncements about how artificial intelligence will change the economy is justified. Nonetheless, there are good reasons to take seriously the growing potential of AI—systems that exhibit intelligent behavior, such as learning, reasoning, and problem-solving—to transform the economy, especially given the astonishing technical advances of the past year.

AI may affect society in a number of areas besides the economy—including national security, politics, and culture. But in this article, we focus on the implications of AI on three broad areas of macroeconomic interest: productivity growth, the labor market, and industrial concentration. AI does not have a predetermined future. It can develop in very different directions. The particular future that emerges will be a consequence of many things, including technological and policy decisions made today. For each area, we present a fork in the road: two paths that lead to very different futures for AI and the economy. In each case, the bad future is the path of least resistance. Getting to the better future will require good policy—including

- Creative policy experiments
- A set of positive goals for what society wants from AI, not just negative outcomes to be avoided
- Understanding that the technological possibilities of AI are deeply uncertain and rapidly evolving and that society must be flexible in evolving with them

First fork: Productivity growth

The first road concerns the future of economic growth—which is largely the future of productivity growth. The US economy has been stuck with disturbingly low productivity growth for most of the past 50 years, except for a brief resurgence in the late 1990s and early 2000s (Brynjolfsson, Syverson, and Chad 2019). Most advanced economies now have the same problem of low productivity growth. More than any other factor, productivity—output per unit of input—determines the wealth of nations and the living standards of their people. With higher productivity, such problems as budget deficits, poverty reduction, health care, and the environment become far more manageable. Boosting productivity growth may be the globe's most fundamental economic challenge.

Low-productivity future

On one path of the productivity fork, AI's impact is limited. Despite the rapidly improving technical capabilities of AI, its adoption by businesses may continue to be slow and confined to large firms (Zolas and others 2021). The economics of AI may turn out to be of a very narrow labor-saving variety (what Daron Acemoglu and Simon Johnson call a "so-so technology," such as an automated grocery checkout stand), instead of one that enables workers to do something novel or powerful (see "Rebalancing AI" in this issue of F&D). Displaced workers might disproportionately end up in even less productive and less dynamic jobs, further muting any aggregate benefit to the long-term productivity growth rate of the economy.

Like so many of Silicon Valley's recent technological enthusiasms (3D printers, self-driving cars, virtual reality), AI may also end up being less promising or less ready to bring to market than initially hoped. Any real economic gains, even modest ones, may show up in the data many decades after the first moments of technological promise, as has often been the pattern. The famous paradox identified by economist Robert Solow in 1987—"You can see the computer age everywhere but the productivity statistics"—may become more extreme, as everyone seems to have an AI chatbot that amazes their friends, but businesses do not seem more productive for their increased use of AI. Firms may further blunt any economic benefits from AI by failing to figure out the organizational and managerial changes they need to best leverage it.

And, as in the case of self-driving cars, the technological challenges of going from an exciting proof of concept to a highly reliable product may be further compounded by a legal regime that was not designed to accommodate this new technology and may seriously hinder its development. In the case of AI, there is tremendous uncertainty over what current laws concerning intellectual property imply when models are trained on millions of data points that may include the

"The path that leads to a worse future is the one of least resistance and results in low productivity growth, higher income inequality, and higher industrial concentration."

protected intellectual property of others. Intellectual property law may eventually respond by creating something analogous to a "patent thicket" that effectively prevents models from being trained on data to which the developers do not have clear rights. At the same time, the wrong choices could undermine the incentives of creative professionals to produce more of the novel content that powers machine learning systems.

In addition, national regulators, driven by any number of concerns, may impose strict regulations that slow the speed of AI development and dissemination. They may even be urged on by the early developers of AI who are eager to protect their lead. Moreover, some countries, businesses, and other organizations may totally ban AI.

High-productivity future

But there is an alternate scenario in which AI leads to a higher-productivity-growth future. AI might be applied to a substantial share of the tasks done by most workers (Eloundou and others 2023) and massively boost productivity in those tasks. In this future, AI lives up to its promise of being the most radical technological breakthrough in many decades. Moreover, it ends up complementing workers—freeing them to spend more time on nonroutine, creative, and inventive tasks rather than just replacing them. AI captures and embodies the tacit knowledge (acquired through experience but hard to articulate) of individuals and organizations by drawing on vast amounts of newly digitized data. As a result, more workers can spend more time working on novel problems, and a growing share of the labor force increasingly comes to resemble a society of research scientists and innovators. The result is an economy not simply at a higher level of productivity, but at a permanently higher growth rate.

In this future, the successful integration of AI with robots also means that much more of the economy is amenable to AI-related progress. And AI enables society not just to do better the things it already does but to do things and envision things previously unimaginable. AI-backed research in medicine enables radical advances in knowledge of human biology and drug design. AI becomes capable of helping the engine of creativity and scientific discovery itself—math, science, further AI development—a kind of recursive self-improvement that was once just a science fiction thought experiment.

Second fork: Income inequality

The increase in income inequality between individual workers over the past 40 years is a major concern. A large body of empirical research in labor economics suggests that computers and other forms of information technology may have contributed to income inequality by automating away routine middle-income jobs, which has polarized the labor force into high-income and low-income workers. Although the CEO and the janitor remain, computers have replaced some of the middle tier of office workers (Autor, Levy, and Murnane 2003). We consider two scenarios for AI's effect on inequality.

Higher-inequality future

In the first scenario, AI leads to higher income inequality. Technologists and managers design and implement AI to substitute directly for many kinds of human labor, driving down the wages of many workers. To make matters worse, generative AI starts to produce words, images, and sounds, tasks formerly thought of as nonroutine and even creative—enabling machines to interact with customers and create the content for a marketing campaign. The number of jobs under threat from AI competition eventually grows much larger. Entire industries are upended and increasingly replaced (a threat to labor perhaps foreshadowed by the recent strikes of screenwriters and actors in the United States, who demanded that studios restrict their use of AI).

This is not a future of mass unemployment. But in this higher-inequality future, as AI substitutes for high- or decently paying jobs, more workers are relegated to low-paying service jobs—such as hospital orderlies, nannies, and doormen—where some human presence is intrinsically valued and the pay is so low that businesses cannot justify the cost of a big technological investment to replace them. The final bastion of purely human labor may be these types of jobs with a physical dimension. Income inequality increases in this scenario as the labor market is further polarized into a small, high-skilled elite and a large underclass of poorly paid service workers.

Lower-inequality future

In the second scenario, however, AI leads to lower income inequality because its main impact on the workforce is to help the least experienced or least knowledgeable workers be better at their jobs. Software coders, for instance, now benefit from the assistance of AI models, such as Copilot, which effectively

draw on coding best practices from many other workers. An inexperienced or subpar coder using Copilot becomes more comparable to a very good coder, even when both have access to the same AI. A study of 5,000 workers who do complex customer assistance jobs at a call center found that among workers who were given the support of an AI assistant, the least skilled or newest workers showed the greatest productivity gains (Brynjolfsson, Li, and Raymond 2023). If employers shared these gains with workers, distribution of income would become more equal.

In addition to creating a future of lower income inequality, AI may help labor in another more subtle, but profound, sense. If AI is a substitute for the most routine and formulaic kinds of tasks, then by taking tedious routine work off human hands, AI may complement genuinely creative and interesting tasks, improving the basic psychological experience of work, as well as the quality of output. Indeed, the call center study found not only productivity gains, but reduced worker turnover and increased customer satisfaction for those using the AI assistant.

Third fork: Industrial concentration

Since the early 1980s, industrial concentration—which measures the collective market share of the largest firms in a sector—has risen dramatically in the United States and many other advanced economies. These large superstar firms are often much more capital-intensive and technologically sophisticated than their smaller counterparts.

There are again two divergent scenarios for the impact of AI.

Higher-concentration future

In the first scenario, industrial concentration increases, and only the largest firms intensively use AI in their core business. AI enables these firms to become more productive, profitable, and larger than their competitors. AI models become ever more expensive to develop, in terms of raw computational power—a massive up-front cost that only the largest firms can afford—in addition to requiring training on massive datasets, which very large firms already have from their many customers and small firms do not. Moreover, after an AI model is trained and created, it can be expensive to operate. For example, the GPT-4 model cost more than $100 million to train during its initial development and requires about $700,000 a day to run. The typical cost of developing a large AI model may soon be in the billions of dollars. Executives at the leading AI firms predict that the scaling laws that show a strong relationship between increases in training costs and improved performance will hold for the foreseeable future, giving an advantage to the companies with access to the biggest budgets and the biggest datasets.

It may be, then, that only the largest firms and their business partners develop proprietary AI—as firms such as Alphabet, Microsoft, and OpenAI have already done and smaller firms have not. The large firms then get larger.

More subtly, but perhaps more important, even in a world in which proprietary AI does not require a large fixed cost that only the largest firms can afford, AI might still disproportionately benefit the largest firms, by helping them better internally coordinate their complex business operations—of a kind that smaller and simpler firms do not have. The "visible hand" of top executives managing resources inside the largest firms, now backed by AI, allows the firm to become even more efficient, challenging the Hayekian advantages of small firms' local knowledge in a decentralized market.

Lower-concentration future

In the lower-industrial-concentration future, however, open-source AI models (such as Meta's LLaMA or Berkeley's Koala) become widely available. A combination of for-profit companies, nonprofits, academics, and individual coders creates a vibrant open-source AI ecosystem that enables broad access to developed AI models. This gives small businesses access to industry-leading production technologies they could never have had before.

Much of this was foreshadowed in an internal memo leaked from Google in May 2023, in which a researcher said that "open-source models are faster, more customizable, more private, and pound-for-pound more capable" than proprietary models. The researcher said that processes in small open-source models can be quickly repeated by many people and end up better than large private models that are slowly iterated by a single team and that open-source models can be trained more cheaply. In the Google researcher's view, open-source AI may end up dominating the expensive proprietary models.

It may also be that AI encourages the kind of broad, decentralized innovation that better flourishes across many small firms than within one large firm. The boundaries of the firm are the outcome of a series of tradeoffs; a world in which more AI-backed innovators need

the residual control rights to their work might be one in which more innovators decide they would rather be owners of small firms than be employees of large ones.

The result is that the long rise in industrial concentration starts to run aground, because some nimble smaller businesses close or even reverse the technology gap with their larger counterparts and win back more market share.

Toward a policy agenda

For each of the forks in the road, the path that leads to a worse future is the one of least resistance and results in low productivity growth, higher income inequality, and higher industrial concentration. Getting to the good path of the fork will require hard work—smart policy interventions that help shape the future of technology and the economy.

It is also important to appreciate a broader point about policy. Much of the discourse around AI regulation now takes place along a kind of hydraulic model: should we have more AI or less AI—or even ban AI. This discussion happens when AI is perceived as somewhat of a fixed thing, with a predetermined future. AI can come fast or slow. There can be more or less of it, but basically it is what it is.

However, if policymakers understand that AI can develop in different directions, the discourse will be framed differently. How can policies encourage the types of AI that complement human labor instead of imitating and replacing it? What choices will encourage the development of AI that firms of all sizes can access, instead of just the largest ones? What kind of open-source ecosystem might that require, and how do policymakers support it? How should AI labs approach model development, and how should firms approach AI implementation? How does society get an AI that unleashes radical innovation, instead of marginal tweaks to existing goods, services, and systems?

Many different actors have power to affect the direction of the AI future. Major corporations will have to make important decisions about how they choose to integrate AI into their workforce. The largest of these companies will also develop in-house AI. AI/computer science labs at universities will also develop AI models, some of which they will make open-source. Federal legislators and regulators will have a large impact, as might more local ones. Voters have a voice. Labor unions must figure out what kind of relationship they want with AI and what their demands will be.

Although we have sketched a number of possible futures for AI, we want to emphasize not only how deeply unpredictable the future of this technology is but also the agency society has in actively and collectively determining which AI future emerges.

We have raised more questions than we have answered, which reflects, in part, the nascent stage of AI adoption and impact. But it also reflects a deeper imbalance between research efforts advancing the frontier of the technology and the more limited research aimed at understanding its economic and social consequences.

This imbalance was of less significance when the technology had limited macroeconomic consequences. But today, when the effects of AI on society are likely to be measured in trillions of dollars, far greater investment should be made in research on the economics of AI. Society needs innovations in economic and policy understanding that match the scale and scope of the breakthroughs in AI itself. Reorienting research priorities and developing a smart policy agenda can help society move toward a future of both sustained and inclusive economic growth. F&D

"Society needs innovations in economic and policy understanding that match the scale and scope of the breakthroughs in AI itself."

ERIK BRYNJOLFSSON *is the Jerry Yang and Akiko Yamazaki Professor at the Stanford Institute for Human-Centered AI, where he directs the Stanford Digital Economy Lab.* **GABRIEL UNGER** *is a postdoctoral fellow at the Stanford Digital Economy Lab.*

REFERENCES

Autor, David, Frank Levy, and Richard Murnane. 2003. "The Skill Content of Recent Technological Change." *Quarterly Journal of Economics* 118 (4): 1279–333.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2019. "Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics." In *The Economics of Artificial Intelligence: An Agenda*, edited by Ajay Agrawal, Joshua Gans, and Avi Goldfarb. Chicago: University of Chicago Press.

Brynjolfsson, Erik, Danielle Li, and Lindsay Raymond. 2023. "Generative AI at Work." NBER Working Paper 31161, National Bureau of Economic Research, Cambridge, MA.

Eloundou, Tyna, Sam Manning, Panels Mishkin, and Daniel Rock. 2023. "GPTs Are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models." *arXiv preprint arXiv:2303.10130*.

Zolas, Nicholas, Zachary Kroff, Erik Brynjolfsson, Kristina McElheran, David N. Beede, Cathy Buffington, Nathan Goldschlag, Lucia Foster, and Emin Dinlersoz. 2021. "Advanced Technologies Adoption and Use by U.S. Firms: Evidence from the Annual Business Survey." NBER Working Paper 28290, National Bureau of Economic Research, Cambridge, MA. https://www.nber.org/papers/w28290.