

Machines of mind: The case for an AI-powered productivity boom

Martin Neil Baily, Erik Brynjolfsson, and Anton Korinek

May 10, 2023

Large language models such as ChatGPT are emerging as powerful tools that not only make workers more productive but also increase the rate of innovation, laying the foundation for a significant acceleration in economic growth. As a general purpose technology, AI will impact a wide array of industries, prompting investments in new skills, transforming business processes, and altering the nature of work. However, official statistics will only partially capture the boost in productivity because the output of knowledge workers is difficult to measure. The rapid advances can have great benefits but may also lead to significant risks, so it is crucial to ensure that we steer progress in a direction that benefits all of society.

On a recent Friday morning, one of us sat down in his favorite coffee shop to work on a new research paper regarding how AI will affect the labor market. To begin, he pulled up [ChatGPT ↗](#), a generative AI tool. After entering a few plain-English prompts, the system was able to provide a suitable economic model, draft code to run the model, and produce potential titles for the work. By the end of the morning, he had achieved a week's worth of progress on his research.

We expect millions of knowledge workers, ranging from doctors and lawyers to managers and salespeople to experience similar ground-breaking shifts in their productivity within a few years, if not sooner.

The potential of the most recent generation of AI systems is illustrated vividly by the viral uptake of ChatGPT, a large language model (LLM) that captured public attention by its ability to generate coherent and contextually appropriate text. This is not an innovation that is languishing in the basement. Its capabilities have already captivated hundreds of millions of users.

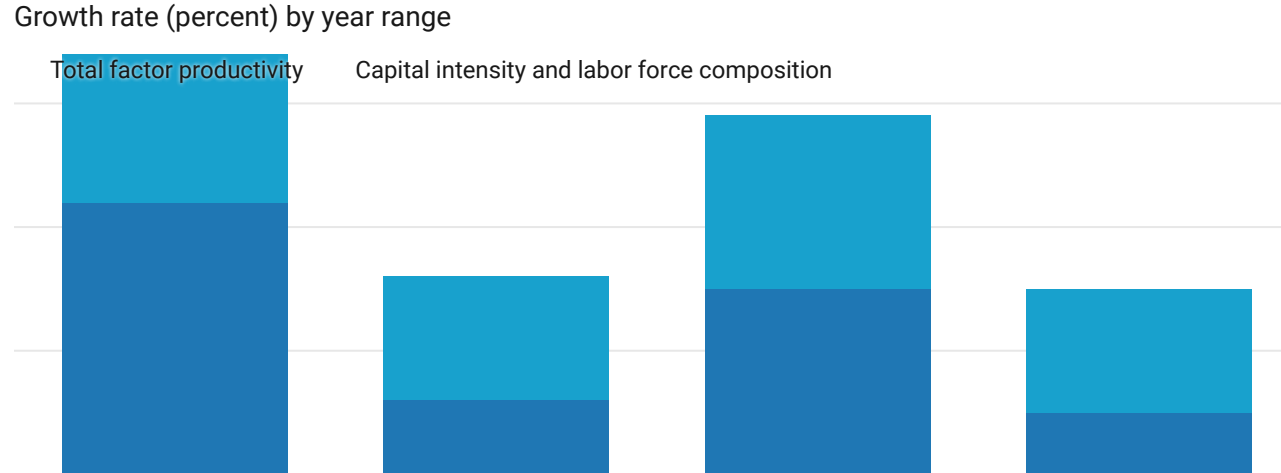
Other LLMs that were recently rolled out publicly include Google's [Bard ↗](#) and Anthropic's [Claude ↗](#). But generative AI is not limited to text: in recent years, we have also seen generative AI systems that can create images, such as [Midjourney ↗](#), [Stable Diffusion ↗](#) or [DALL-E ↗](#), and more recently multi-modal

systems that combine text, images, video, audio and even [robotic functions](#) ↗. These technologies are [foundation models](#) ↗, which are vast systems based on deep neural networks that have been trained on massive amounts of data and can then be adapted to perform a wide range of different tasks. Because information and knowledge work dominates the US economy, these machines of the mind will dramatically boost overall productivity.

The power of productivity growth

The primary determinant of our long-term prosperity and welfare is the rate of productivity growth: the amount of output created per hour worked. This holds even though changes in productivity are not immediately felt by everyone and, in the short run, workers' perceptions of the economy are dominated by the business cycle. From World War II until the early 1970s, labor productivity grew at over 3% a year, more than doubling over the period, ushering in an era of prosperity for most Americans. In the early 1970s productivity growth slowed dramatically, rebounding in the 1990s, only to slow again since the early 2000s.

Figure 1. The components of labor productivity growth 1948 to 2022.



2005-2022

Source: Bureau of Labor Statistics, Productivity Database • [Get the data](#)

Figure 1 illustrates the story. It decomposes the overall growth in labor productivity into two components: total factor productivity (which is a measure of the impact of technology) and the contribution of the labor composition and capital intensity. The figure illustrates that the key driver of changes in labor productivity is changes total factor productivity (TFP). There are many reasons for America's recent economic struggles, but slow TFP growth is a key cause, slowly eating away at the country's prosperity, making it harder to fight inflation, eroding workers' wages and worsening budget deficits.

The generally slow pace of economic growth, together with the outsized profits of tech companies, has resulted in skepticism about the benefits of digital technologies for the broad economy. However, for about 10 years starting in the 1990s there was a surge in productivity growth, as shown in Figure 1, driven primarily by a huge [wave of investment in computers and communications](#), which in turn drove business transformations. Even though there was a stock market bubble as well as significant reallocation of labor and resources, workers were generally better off. Furthermore, the federal budget [was balanced from 1998 to 2001](#)—a double win. Digital technology can drive broad economic growth, and it happened less than thirty years ago.

Early estimates of AI's productivity effects

The recent advances in generative AI have been driven by progress in software, hardware, data collection, and growing amounts of investment in cutting-edge models. [Sevilla et al. \(2022\)](#) observe that the amount of compute (computing power) used to train cutting-edge AI systems has been doubling every six months over the past decade. The capabilities of generative AI systems have grown in tandem, allowing them to perform many tasks that used to be reserved for cognitive workers, such as writing well-crafted sentences, creating computer code, summarizing articles, brainstorming ideas, organizing plans, translating other languages, writing complex emails, and much more.

Generative AI has broad applications that will impact a wide range of workers, occupations, and activities. Unlike most advances in automation in the past, it is a machine of the mind affecting cognitive work. As noted in a recent research paper ([Eloundou et al., 2023](#)) ↗, LLMs could affect 80% of the US workforce in some form.

There is an emerging literature that estimates the productivity effects of AI on specific occupations or tasks. [Kalliamvakou \(2022\)](#) ↗ finds that software engineers can code up to twice as fast using a tool called Codex, based on the previous version of the large language model GPT-3. That's a transformative effect. [Noy and Zhang \(2023\)](#) ↗ find that many writing tasks can also be completed twice as fast and [Korinek \(2023\)](#) ↗ estimates, based on 25 use cases for language models, that economists can be 10-20% more productive using large language models.

But can these gains in specific tasks translate into significant gains in a real-world setting? The answer appears to be yes. [Brynjolfsson, Li, and Raymond \(2023\)](#) ↗ show that call center operators became 14% more productive when they used the technology, with the gains of over 30% for the least experienced workers. What's more, customer sentiment was higher when interacting with operators using generative AI as an aid, and perhaps as a result, employee attrition was lower. The system appears to create value by capturing and conveying some of the tacit organizational knowledge about how to solve problems and please customers that previously was learned only via on-the-job experience.

Criticism of large language models as merely “stochastic parrots” is misplaced. Most cognitive work involves drawing on past knowledge and experience and applying it to the problem at hand. It is true that generative AI programs are prone to certain types of mistakes, but the form of these mistakes is predictable. For example, language models tend to engage in [“hallucinations,”](#) i.e., to make up facts and references. As a result, they clearly require human oversight. However, their economic value depends not on whether they are flawless, but on whether they can be used productively. By that criterion, they are already poised to have a massive impact. Moreover, the accuracy of generative AI models continues to improve rapidly.

Quantifying the productivity effects

A [recent report by Goldman Sachs](#) suggests that generative AI could raise global GDP by 7%, a truly significant effect for any single technology. Based on our analysis of a variety of use cases and the share of the workforce doing mainly cognitive work, this estimate strikes us as being reasonable, though there remains great uncertainty about the ultimate productivity and growth effects of AI.

It is useful to rigorously break down the channels through which we expect generative AI to produce growth in productivity, output, and ultimately in social welfare in a model.


The first channel is the increased efficiency of output production. By making cognitive workers engaged in production more efficient, the level of output increases. Economic theory tells us that, in competitive markets, the effect of a productivity boost in a given sector on aggregate productivity and output is equal to the size of the productivity boost multiplied by the size of the sector ([Hulten's theorem](#)). For instance, if generative AI makes cognitive workers on average 30% more productive over a decade or two and cognitive work makes up about 60% of all value added in the economy (as measured by the wage bill attributable to cognitive tasks), this amounts to a 18% increase in aggregate productivity and output, spread out over those years.

The second, and ultimately more important, channel is the acceleration of innovation and thus future productivity growth. Cognitive workers not only produce current output but also invent new things, engage in discoveries, and generate the technological progress that boosts future productivity. This includes R&D—what scientists do—and perhaps more importantly, the process of rolling out new innovations into production activities throughout the economy—what managers do. If cognitive workers are more efficient, they will accelerate technological progress and thereby boost the rate of productivity growth—in perpetuity. For example, if productivity growth was 2% and the cognitive labor that underpins productivity growth is 20% more productive, this would raise the growth rate of productivity by 20% to 2.4%. In a given year, such a change is barely noticeable and is usually swamped by cyclical fluctuations.

But productivity growth compounds. After a decade, the described tiny increase in productivity growth would leave the economy 5% larger, and the growth would compound further every year thereafter. What's more, if the acceleration applied to the growth rate of the growth rate (for instance if [one of the applications of AI was to improving AI itself ↗](#)), then of course, growth would accelerate even more over time.

Figure 2. Possible Growth Trajectories

% growth by year



Source: Authors' simulations. • [Get the data](#)

Figure 2 schematically illustrates the effects of the two channels of productivity growth over a twenty year horizon. The baseline follows the [current projection of the Congressional Budget Office \(CBO\) of 1.5% productivity growth ↗](#), giving rise to a total of 33% productivity growth over

20 years. The projection labeled "Level" assumes that generative AI raises the level of productivity and output by an additional 18% over ten years, as suggested by the illustrative numbers we discussed for the first channel. After ten years, growth reverts to the baseline rate. The third projection labeled "Level+Growth" additionally includes a one percentage point boost in the rate of growth over the baseline rate, resulting from the additional innovation triggered by generative AI. At first, the resulting growth trajectory is barely distinguishable from the "Level" projection, but through the power of compounding, the effects grow bigger over time, leading to a near doubling of output after 20 years, far greater than the baseline projection.

Barriers and drivers of adoption

For the productivity gains to materialize, advances in AI have to disseminate throughout the economy. Traditionally, this has always taken time, so we would not expect potential productivity gains to show up immediately. The advances need to be taken up and rolled out by businesses and organizations that employ cognitive labor throughout the economy, including small and medium-sized businesses, some of which may be slow to realize the potential of adapting advanced new technologies or may lack the required skills to use them well. For example, the Goldman report assumes it takes 10 years for the gains to fully materialize.

The “productivity J-curve” ([Brynjolfsson et al., 2021](#)) ⁷ describes how new technologies, especially general purpose technologies, deliver productivity gains only after a period of investment in complementary intangible goods, such as business processes and new skills. In fact, this can temporarily even drag down measured productivity. As a result, earlier general purpose technologies like electricity and the first wave of computers took decades to have a significant effect on productivity. Additional barriers to adoption and rollout include concerns about job losses and institutional inertia and regulation, in areas from the medicine to finance and law.

However, in the case of generative AI there are also factors that can mitigate these barriers, or even accelerate adoption. First, in contrast to physical automation, one benefit of cognitive automation is that it can often be rolled out quickly via software. This is particularly true now that a ubiquitous digital infrastructure is available: the Internet. ChatGPT famously was the most rapid product launch in history—it gained [100 million users in just two months](#) ⁷—because it was accessible to anyone with an internet connection and did not require any hardware investment on the users’ side.

Both Microsoft and Google are in the process of rolling out Generative AI tools as part of their search engines and office suites, offering access to generative AI to a large fraction of the cognitive workforce in advanced countries who regularly use these tools. Furthermore, application programming interfaces (APIs) are increasingly available to enable seamless

modularization and connectivity between systems, and a marketplace for plug-ins and extensions is rapidly growing, making it much easier to add functionality. Finally, in contrast to other technologies, users of generative AI can interact with the technology in natural language rather than special codes or commands, making it easier to learn and adopt these tools.

These reasons for optimism suggest that the rollout of these new technologies may be faster than in the past. Still, the importance of training to make optimal use of these tools cannot be overstated.

Problems of measurement – silent productivity growth

The most common measure of productivity, non-farm business productivity, is quite adept at capturing increases in [productivity in the industrial sector](#) where inputs and outputs are tangible and easy to account for. However, productivity of cognitive labor is harder to measure. Statisticians who compile GDP and productivity statistics sometimes resort to valuing the output of cognitive activity simply by assuming it is proportional to the quantity of labor input being used to produce it, which of course eliminates any scope for productivity growth.

For example, generative AI enables economists to write more thought pieces and provide deeper analyses of the economy than before, yet this output would not directly show up in GDP statistics. Readers may feel that they have access to better and deeper economic analyses (contributing to channel 1 above). Moreover, the analyses may also play a part in enabling business leaders and policymakers to better harness the positive productivity effects of generative AI (contributing to channel 2 above). Neither of these positive productivity effects of such work would be directly captured in official GDP or productivity statistics, yet the benefits of economists' productivity gains would still lead to greater social welfare.

The same holds true for many other cognitive workers throughout the economy. This may give rise to significant under-measurement or "silent productivity growth."

Productivity growth, labor markets, and income distribution

A bigger pie does not automatically mean everyone benefits evenly, or at all. The productivity effects of generative AI are likely to go hand in hand with significant disruption in the job market as many workers may see downward wage pressures. For example, the Eloundou et al. paper cited earlier predicts that up to 49% of the workforce could *eventually* have half or more of their

job tasks performed by AI. Will the demand for these tasks increase enough to compensate for such efficiency gains? Will the workers find other tasks to do? The answers are far from certain. In past technological transformations, workers who lost their jobs could transition to new jobs, and on average pay increased. However, given the scale of the impending disruption and the labor-saving nature of it, it remains to be seen whether this will be the case in the age of generative AI.

Moreover, the current wave of cognitive automation marks a change from most earlier waves of automation, which focused on physical jobs or routine cognitive tasks. Now, creative and unstructured cognitive jobs are also being impacted. Instead of the lowest paid workers bearing the brunt of the disruption, now many of the highest-paying occupations will be affected. These workers may find the disruption to be quite unexpected. If their skills are general, they may find it easier to adjust to displacement than blue-collar workers. However, if they have acquired a significant amount of human capital that becomes obsolete, they may experience much larger income losses than blue-collar workers who were displaced by previous rounds of automation.

The idea of jobs created versus jobs displaced is the most tangible manifestation of job market disruption for lay people. Job losses are indeed a significant social concern, and we need policies to facilitate adjustment. However, as economists, we note that the key factor in determining the

influence of new technologies on the labor market is ultimately their effect on labor demand. Counting how many jobs are created versus how many are destroyed misses that employment is determined as the equilibrium of labor demand and labor supply. Labor supply is quite inelastic, reflecting that most working-age people want to or have to work independently of whether their incomes go up or down. Workers who lose their jobs as a result of changing technology will seek alternative employment. And, to the extent that changing technology raises productivity, this will increase national income and spur the demand for labor. Over the long run, the labor market can be expected to equilibrate, meaning that the supply of jobs, the demand for jobs and the level of wages will adjust to maintain full employment. This is evidenced by the fact that the unemployment rate in the United States has remained [consistently low in the postwar period ↗](#) (with help from monetary and fiscal policy to recover from recessions). Job destruction has always been offset by job creation. Instead, the effects of automation and augmentation tend to be reflected in wages and income.

The effect of generative AI on labor demand depends on whether the [systems complement or substitute for labor ↗](#). Substitution occurs when AI models automate most or all tasks of certain jobs, while complementing occurs if they automate small parts of certain jobs, leaving humans indispensable. Additionally, AI systems can be complementary to human labor if they enable new tasks or increase quality.

As companies invest more in generative AI, they often have choices about whether to emphasize substitution or complementarity. For example, if call centers can use AI to complement human operators, or, as AI improves, they may restructure their processes to have the systems address more and more queries without human operators being involved. At the same time, higher productivity growth across the economy may make the overall effects more complementary by increasing overall labor demand and may mitigate the disruption.

In recent decades, there have been three main forces impacting income distribution. First, there has been an overall shift of income away from wages and towards corporate capital. Second, there has been an [increase in the return to the skills ↗](#) that are valued by companies (reflected in part by higher returns to education). Third, there has been a shift caused by [increased foreign competition ↗](#).

It is hard to predict how generative AI will impact this mix. A positive interpretation is that workers who currently struggle with aspects of math and writing will become more productive with the help of these new tools and will be able to take better-paid jobs with the help of the new technology. A negative interpretation is that companies will use the technology to eliminate or de-skill more and more positions pushing a larger fraction of the workforce into unfulfilling jobs, raising the share of profits in income and,

perhaps, increasing the demand for the most elite members of the workforce.

No doubt technological progress will not stop with the current wave of generative AI. Instead, we can expect even more dramatic advances in AI, bringing the technology closer to what is called artificial general intelligence (AGI). This [will lead to even more radical transformations of life and work ↗](#). The scarcity of human labor has been [a double-edged sword throughout our history ↗](#): on the one hand, it has held back economic growth because greater production would require more labor; on the other hand, it has been highly beneficial for income distribution since wages represent the market value of scarce labor. If labor can be replaced by machines across a wide range of tasks in the future, both points may no longer hold, and we may experience an AI-powered growth take-off at the same time as that the value of labor declines. This would present a [significant challenge for our society \(https://www.brookings.edu/research/preparing-for-the-non-existent-future-of-work/\) .](#) Moreover, AGI may also impose large risks on humanity [if not aligned with human objectives \(https://www.brookings.edu/research/aligned-with-whom-direct-and-social-goals-for-ai-systems/\) .](#)

Conclusion

Large language models and other forms of generative AI are still at an early stage, making it difficult to predict with great confidence the exact productivity effects they will have. Yet as we have argued, we expect that generative AI will have tremendous positive productivity effects, both by increasing the level of productivity and accelerating future productivity growth.

For policymakers, the goal should be to allow for the positive productivity gains while mitigating the risks and downsides of ever-more powerful AI. Faster productivity growth is an elixir that can solve or mitigate many of our society's challenges, from raising living standards and addressing poverty to providing healthcare for all and strengthening our defenses. Indeed, it will be nearly impossible to fix some of our budgetary challenges, including the growing deficits, without sufficiently stronger growth.

AI-powered productivity growth will also create challenges. There may be a need for updating social programs and tax policy to soften the welfare costs of labor market disruptions and ensure that the benefits of AI give rise to shared prosperity rather than concentration of wealth. Other harms will also need to be addressed, including the amplification of misinformation and polarization, potentially destabilizing our democracy, and the creation of new biological and other weapons that could injure or kill untold numbers of people.

Therefore, we cannot let the capabilities of AI outstrip our understanding of their potential impacts. Economists and other social scientists will need to accelerate their work on AI's impacts to keep up with our colleagues in AI research who are rapidly advancing the technologies. If we do that, we are optimistic our society can harness the productivity benefits and growth acceleration delivered by artificial intelligence to substantially advance human welfare in the coming years.

The authors used GPT4 for writing assistance in producing this text but assume full responsibility for its content and accuracy.


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 \(https://www.brookings.edu/about-us/annual-report/\)](https://www.brookings.edu/about-us/annual-report/). The findings, interpretations, and conclusions in this report are solely those of its author(s) and are not influenced by any donation.

AUTHORS




Martin Neil Baily Senior Fellow Emeritus - Economic Studies, Center on Regulation and Markets



Erik Brynjolfsson Director - Stanford Digital Economy Lab, Jerry Yang and Akiko Yamazaki Professor and Senior Fellow - Stanford Institute for Human Centered AI  @erikbryn



Anton Korinek Nonresident Senior Fellow - Economic Studies, Center on Regulation and Markets  @akorinek

The Brookings Institution is committed to quality, independence, and impact. We are supported by a [diverse array of funders \(/about-us/annual-report/\)](/about-us/annual-report/). In line with our [values and policies \(/about-us/research-independence-and-integrity-policies/\)](/about-us/research-independence-and-integrity-policies/), each Brookings publication represents the sole views of its author(s).