AI and Growth: Where Do We Stand?*

Philippe Aghion Simon Bunel

June 2024

Abstract

In this note we use two alternative approaches to estimate the macroeconomic impact of artificial intelligence (AI) on productivity growth over the next decade. The first approach exploits the parallel between the AI revolution and past technological revolutions. The second approach follows Acemoglu (2024) and the task-based framework, which we revisit using our own reading of the existing empirical literature on the various component of the task-based formula. Based on the first approach, we estimate that the AI revolution should increase aggregate productivity growth by between 0.8 and 1.3pp per year over the next decade. Using the second approach but with our own reading of the recent empirical literature on the various components of the task-based formula, we obtain a median estimate of 0.68pp additional annual total factor productivity (TFP) growth. Our estimates do not take into account the fact that AI automates tasks not only in the production of goods and services, our focus in this note, but also in the production of ideas.

*Addresses - Aghion: Collège de France, INSEAD et London School of Economics; Bunel: Banque de France et École Normale Supérieure.

1 Introduction

The impact of IA on potential productivity growth can occur through two distinct channels: (i) increasing productivity, i.e., the rate at which we produce goods and services; and (ii) increasing our ability to generate new ideas and hence new innovations, new products, or new forms of organization.

In this note we concentrate on the first channel, i.e. on the growth effects of AI through automating tasks in the production of goods and services, as it happened with mechanization in agriculture, the invention of the assembly line in industry or, more recently, the digitalization of a significant part of the economy.

A first piece of evidence that speaks to this channel, is the microeconomic study by Brynjolfsson et al. (2023). In that paper, the authors examine the impact of generative AI on the productivity of workers in a U.S. customer service firm. The firm gradually deployed an AI tool to assist employees responsible for responding to customers via online chat by offering automatically generated responses. They show a significant productivity effect of deploying this tool: namely, the productivity of employees who had access to the AI assistant increased by 14% in the first month of use and stabilized at approximately 25% after three months (Figure 1).[1]

How can we move beyond case studies of individual firms to estimate the economywide impact of AI on economic growth? In this note, we consider two alternative approaches to estimate the impact of AI on potential growth over the next decade. The first approach exploits the parallel between the AI revolution and past technological revolutions. The second approach follows Acemoglu (2024) and the task-based framework, which we revisit using information from existing empirical studies.

Based on the first approach, we estimate that the AI revolution should increase aggregate productivity growth by between 0.8 and 1.3pp per year over the next decade. As for the second approach, it leads Acemoglu (2024) to a much smaller extra growth estimate of 0.07 pp per year over the next decade. Using the same task-based formula, but with our own reading of the recent empirical literature regarding each component of that formula, we estimate that the AI revolution should increase aggregate productivity growth by between 0.07pp and 1.24pp, with a median estimate of 0.68pp

[1] These results apply only to a particular type of job within a particular firm. However, two studies (Noy and Zhang, 2023; Dell'Aqua et al., 2023) look at the use of ChatGPT by highly skilled individuals (such as consultants and managers) in the United States and find that their productivity increased by between 25% and 40% for typical tasks in these professions. These findings suggest that productivity gains are observed across a wide range of occupations and skill levels. Positive effects are also experienced at the firm level. In France, an extensive survey by Pôle Emploi (2023) highlights that 72% of employers using AI reported a positive impact on their employees' performance, in particular by reducing tedious tasks (63%) or the risk of error (51%).

additional annual TFP growth. This estimate in turn may be seen as a lower bound to the extent that it does not account for the fact that AI also automates the production of ideas. On the other hand, it does not take into account potential barriers to growth, in particular those associated with the lack of competition in the upstream segments of the AI value chain.

This note is structured as follows. Section 2 presents the first approach, drawing a parallel with previous industrial revolutions. Section 3 reports our estimates using the task-based model, and section 4 discusses the limitations of these approaches, in particular the increased impact of AI through the generation of new ideas and the downward bias due to lack of competition in the AI value chain. Section 5 presents our conclusions.

Figure 1: Effect of adopting generative AI on the productivity of customer service employees

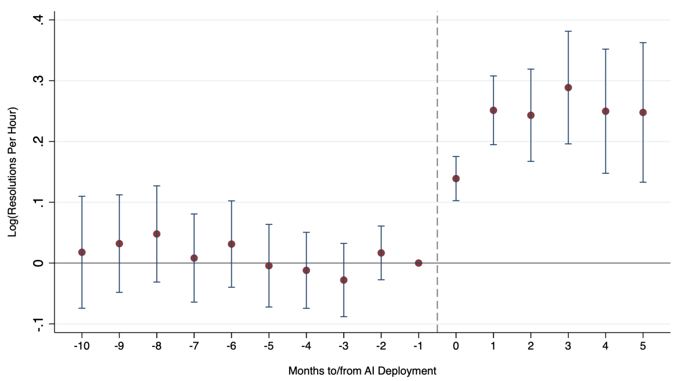

Source: Brynjolfsson et al. (2023)

2 Parallel with previous technological revolutions

In the United States, as in Europe, the productivity gains from electricity did not materialize until about thirty years after the invention of the technology. To understand this time lag, one can look to electricity adoption in industry. In the early 20th century, the use of electricity in manufacturing plants was still limited. Factories maintained an internal organization similar to that of water-powered mills, with a central line shaft turned by hydraulic power. Neither the advent of steam power during the First Industrial Revolution nor the introduction of the dynamo at the beginning of the Second Industrial Revolution led to significant changes in the internal organization of factories. The presence of the line shaft necessitated the placement of similar machines side by side. It was not until the 1910s that the productivity gains associated with electricity were realized, thanks to the introduction of the electrical wire and the miniaturization of electric motors. Each machine then became autonomous, powered by electricity. This innovation eliminated the need for the line shaft and allowed machines to be arranged more efficiently, leading to the development of the assembly line and a subsequent increase in factory productivity (David, 1990).

So too will there be a delay in the impact of AI, as it requires a reorganization of work within firms as well as additional investment. AI was initially developed through the use of deductive "if-then" rules. This symbolic approach, based on reasoning and instructions, was dominant until the 1990s. Although this approach has not been abandoned, a statistical approach to AI, known as machine learning, has gained prominence since the 1990s. Unlike the symbolic approach, in the statistical approach the computer "learns" to identify statistical relationships between data. In this approach, there are no explicit human instructions; instead, the machine is trained to recognize patterns from a set of training data and then applies these patterns to new data to perform a task. Building on the parallel with the introduction of electricity, a thirty-year lag would suggest that the impact of AI on productivity is likely to be felt in the next few years.

This leads us to a second question: what are the expected economic gains? If we assume that the productivity gains enabled by the AI wave of the next decade will be comparable to those of the electricity wave of the 1920s in Europe, then productivity growth would increase by 1.3 percentage points per year starting in 2024 (Figure 2a). If we prefer to use the digital technology wave of the late 1990s and early 2000s in the United States as a point of comparison, the increase in productivity growth would be around 0.8 percentage points per year (Figure 2b). By comparison, France's potential productivity growth is now estimated at 0.5% per year over the medium term.

Figure 2: Effect of previous technological revolutions on productivity growth

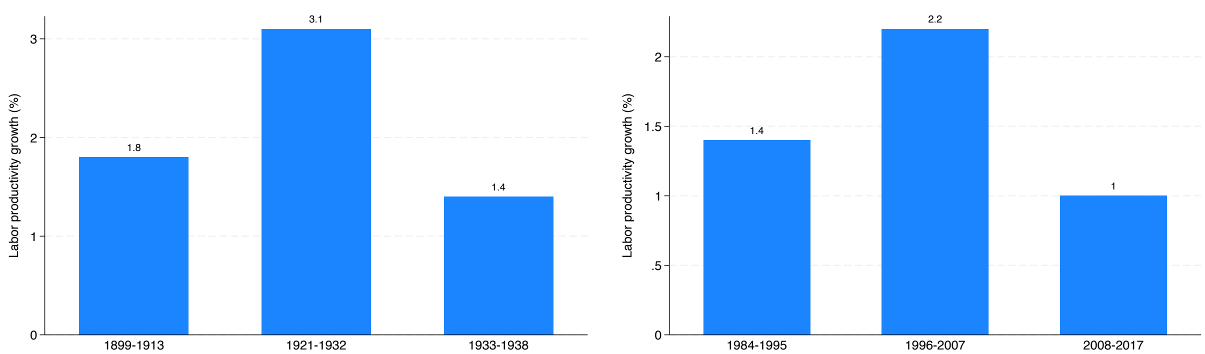

(a) Wave of electricity in Europe

(b) Wave of ICT in the US

Note: In order to avoid the effect of WW1 and the post-war reconstruction of productivity growth, we delete data from the time period 1914-1920. Data from Bergeaud et al. (2016).

This rise in productivity growth would be transitory; once the entire economy has adopted AI, the productivity gains linked to this adoption and subsequent transformations will cease (Figure 3).

Figure 3: Effect of AI adoption in production of goods and services on potential productivity

Source: French Artificial Intelligence Commission (2024)

3 Task-based model approach

In this section, we estimate the effects of AI on TFP by applying the approach set out in Acemoglu (2024), which relies on a theoretical model inspired by Acemoglu and Restrepo (2018). According to this approach, whereas the microeconomic effects of AI are driven by cost savings and productivity improvements at the task level, the macroeconomic effects can be estimated as a function of (i) the fraction of tasks impacted by AI and (ii) the average savings made possible per task:[2]

$$\begin{aligned}\text{TFP gains over 10 years} = {} & \text{GDP share of tasks that are impacted by AI over 10 years} \\ & \text{x Average cost savings in these tasks due to AI}\end{aligned}$$

More precisely, the GDP share of tasks that are impacted by AI over 10 years is the product of (i) the GDP share of tasks impacted by AI and (ii) the share of tasks exposed to AI for which it will be economically profitable to use AI. Meanwhile, the average cost savings realized on the tasks is the

[2]More precisely, this follows from Hulten's theorem, which leads to the following equation $d\ \ln \text{PGF} = \overline{\pi}$ x GDP share of tasks impacted by AI where $\overline{\pi}$ denotes the economy-wide cost savings.

product of (i) the average labor cost savings enabled by AI and (ii) the labor share in value added adjusted for AI exposure. TFP gains over 10 years are then estimated as the combination of 4 terms:

TFP gains over 10 years = GDP share of tasks that are exposed to AI

x Share of exposed tasks for which it will be economically profitable to use AI

x Average costs savings from AI

x AI exposure-adjusted labor share

Estimating each of the 4 terms, Acemoglu (2024) concludes that AI will lead to TFP gains of around 0.7% over the next 10 years, i.e., an increase in annual TFP growth of around 0.07pp. The remaining part of this section discusses the magnitude of these 4 terms, in light of available economics and AI literature.

3.1 GDP share of tasks that are exposed to AI ($ExpAI$)

Numerous papers have been produced on the long-term exposure of tasks to AI. Eloundou et al. (2023) focuses on the impact of large language models (LLMs) - in particular Generative Pre-trained Transformers (GPTs) - on tasks. The authors estimate that a job is exposed to AI if the technology can significantly reduce the time required for certain tasks, including both tasks exposed to the risk of replacement and those with potential for augmentation by AI. They estimate that the GDP share of tasks that are exposed to AI is 19.9%. Following a related approach based on an estimation of the exposure to AI of every task defined by O*NET, Gmyrek et al. (2023) finds 18.5% exposed in advanced countries. Conversely, relying on abilities rather than tasks, Pizzinelli et al. (2023) estimates that this share is much higher, around 60% in the United States, and even reaching 68% in the United Kingdom (Figure 4). Although this study does not provide the share of GDP for tasks exposed to AI but only the share of tasks exposed to AI, the magnitude is likely similar, assuming that exposed tasks are not concentrated within a few specific sectors.

Figure 4: Share of tasks that are exposed to AI

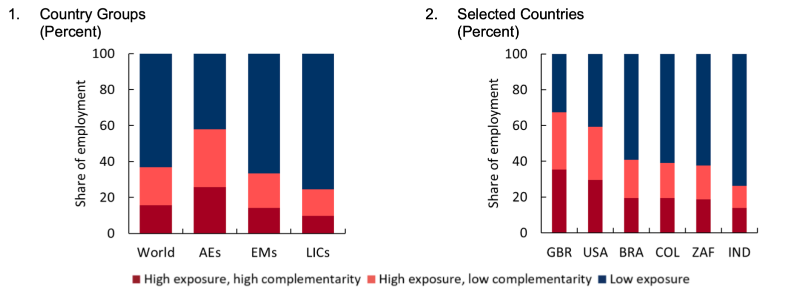

Sources: American Community Survey; Gran Encuesta Integrada de Hogares; India Periodic Labour Force Survey; International Labour Organization; Labour Market Dynamics in South Africa; Pesquisa Nacional por Amostra de Domicílios Contínua; UK Labour Force Survey; and IMF staff calculations.
Note: Country labels use International Organization for Standardization (ISO) country codes. AEs = advanced economies; EMs = emerging market economies; LICs = low-income countries; World = all countries in the sample. Share of employment within each country group is calculated as the working-age-population-weighted average.

Source: Gmyrek et al. (2023)

This leads us to consider that the relevant interval for estimating the share of GDP of tasks exposed to AI is $[\mathbf{0.185}; \mathbf{0.68}]$.

3.2 Share of exposed tasks for which it will be economically profitable to use AI ($ProfitableAI$)

The main limitation of the literature on tasks exposed to AI is that it does not take into account the technical feasibility and economic attractiveness of using AI to replace or augment a job. Determining for which job and when it will be profitable to use AI is a difficult undertaking. Svanberg et al. (2024) studies this share for the specific case of computer vision because this field has more relevant information on cost modeling, and estimates that for U.S. firms, at current costs, the cost/benefit ratio would favor adopting AI technology in only 23% of tasks related to computer vision. However, this figure of 23% suffers from two major shortcomings:

1. It assumes that the proportion of tasks for which it will be profitable to use AI will remain constant over time. This is equivalent to assuming that the cost-saving advances from AI implementation will be zero over the next decade.

2. Svanberg et al. (2024) focuses on computer vision, rather than generative AI in general, which could have a different share of profitable tasks, as it encompasses a much wider technological field.

With regard to the first point, Svanberg et al. (2024) estimates the proportion of tasks that have the potential to incorporate AI that will effectively incorporate the technology. They propose three scenarios with annual cost reductions of 10%, 20%, and 50% respectively (Figure 5).

Figure 5: Share of exposed tasks for which it will be economically profitable to use AI based on different level of cost-saving advances of AI

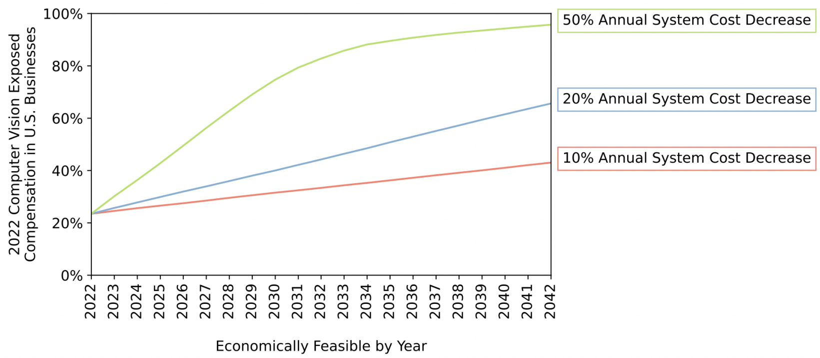

Source: Svanberg et al. (2024)

The 50% annual cost reduction scenario is based on the current rate of change in GPU computing power. In particular, Sevilla et al. (2022) highlights that the computing power of LLMs has doubled every 9 to 10 months since 2016 (Figure 6). In this case, the proportion of tasks for which incorporation of AI would be profitable would increase to approximately 80% within 10 years (Figure 5).

Figure 6: Number of operations per second over time in AI models

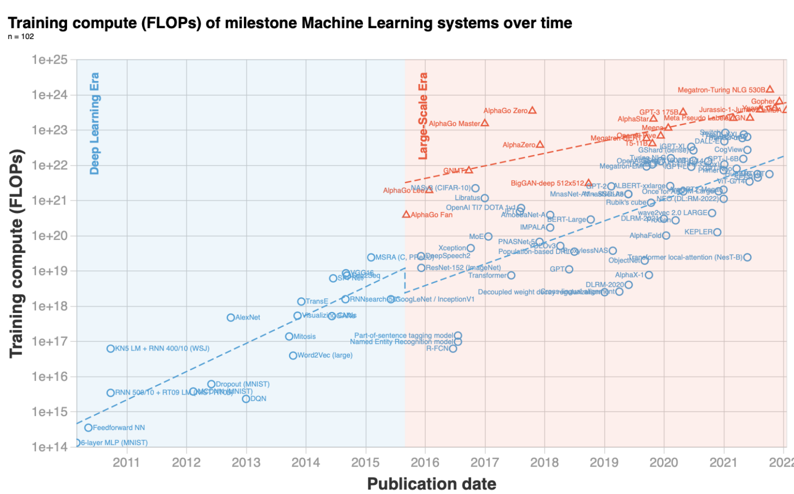

Source: Sevilla et al. (2022)

However, the doubling of GPU capacity does not necessarily translate into a 50% reduction in cost, in particular due to diminishing returns or to energy costs associated with the increase in computing power. Using a methodology based on the cost of computing time rather than performance, Besiroglu and Hobbhahn (2022) estimates a 22% annual decline in computing costs over recent years (Figure 7). At this rate, the proportion of tasks that can incorporate IA will rise to around 50% within 10 years (Figure 5).

Figure 7: Computing power per dollar over time in AI models

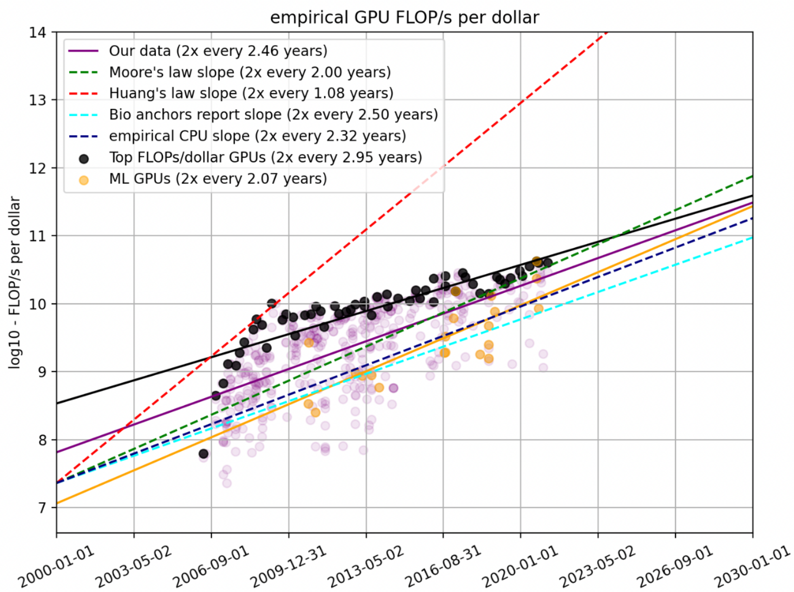

Source: Besiroglu and Hobbhahn (2022)

In respect of the second limitation above, while progress in computer vision has been extremely rapid in recent years, it has not been as spectacular as, for example, recent progress in the ability to understand handwriting or speech (Figure 8). We might therefore expect the developments projected by Svanberg et al. (2024) in the case of computer vision to be faster if we consider the case of generative AI as a whole.

Figure 8: AI capabilities compared to human performance in various fields over time

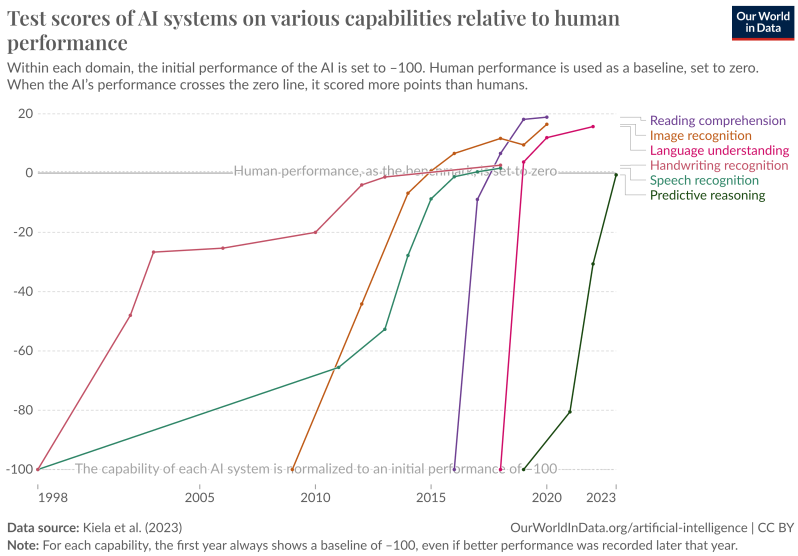

Source: Our World in Data

Without taking this second limitation into account, this leads us to consider that, given the literature available, a relevant interval for estimating the proportion of exposed tasks for which it will be profitable to use AI is $[\mathbf{0.23}; \mathbf{0.8}]$.

3.3 Average costs savings from AI ($LaborCostSavingsAI$)

In order to estimate the cost savings enabled by AI use, we follow Acemoglu (2024) and focus on articles that study the effect of AI on productivity for different occupations, considering that productivity improvements translate into labor cost reductions. Peng et al. (2023) estimates these gains at +55.8% for programmers, and Noy and Zhang (2023) estimates a +40% gain for analysts. For customer service employees, Brynjolfsson et al. (2023) finds a gain of +14% for the first month following the introduction of the AI assistant, reaching +25% for the second month and stabilizing between +25% and +30% in the following three months (Figure 1). Over 10 years, the gains considered for this study thus appear to be in the order of +25%. On the grounds that the analysis framework of Peng et al. (2023) is less relevant because the task evaluated is too finely defined, Acemoglu (2024) looks to the average of the results of the other two studies to derive an effect on worker productivity of +27%. When we take the three studies into account, the long-term productivity effect reaches an average of +40%. This leads us to consider that the relevant interval for estimating the average labor cost savings made possible by AI is $[\mathbf{0.27}; \mathbf{0.4}]$.

3.4 AI exposure-adjusted labor share ($LaborShareAI$)

In the absence of studies providing more precise information on the labor share in value added adjusted for exposure to AI or examining countries other than the United States, we cannot refine the labor share following AI exposure. We therefore retain Acemoglu (2024)'s value of 0.57.

3.5 TFP gains over 10 years

Overall, taking into account all the value intervals for the 4 terms, we infer that the development of AI would lead to an increase in annual productivity growth within a range of [**0,07pp;1,24pp**] over 10 years. This result illustrates, first and foremost, the uncertainty surrounding the quantification of the impact of AI on aggregate productivity growth, depending on the studies from which the estimates are drawn. However, a baseline scenario taking into account (i) the share of tasks exposed to AI in developed countries estimated at 60% (Pizzinelli et al., 2023), (ii) the share of exposed tasks for which it will be profitable to use AI estimated at 50% due to a 22% annual decline in computing costs (Besiroglu and Hobbhahn, 2022) and (iii) productivity gains enabled by AI estimated at 40% based on three benchmark studies (Peng et al., 2023; Noy and Zhang, 2023; Brynjolfsson et al., 2023) leads to an estimated increase in annual productivity growth of **0.68pp** over 10 years, an effect of the same order of magnitude as that presented in the French Artificial Intelligence Commission (2024) report:

$$\text{Annual TFP gains} = \underbrace{ExpAI}_{0.60} \text{ x } \underbrace{ProfitableAI}_{0.50} \text{ x } \underbrace{LaborCostSavingsAI}_{0.40} \text{ x } \underbrace{LaborShareAI}_{0.57} \text{x10} = 0.68$$

4 Discussion

Some will see this estimation as too pessimistic and others as too optimistics. The former will argue that AI can also automate the production of ideas, thereby generating perpetual additional growth. The latter will point to the existence of barriers to growth, in particular the lack of competition in various segments of the AI value chain.

4.1 Automating the production of ideas

AI could automate, or at least facilitate, the generation of new ideas (Aghion et al., 2018). It will thus help us generate new inventions and solve complex problems, as in the case of AlphaFold, which helps find new proteins, or GNoME, which suggests new materials that could be used in vehicles or everyday objects. The impact of AI on science and innovation is difficult to quantify, especially as AI's ability to generate new ideas could face practical difficulties. For instance, it is not enough to identify several million potential new materials; they must still be validated experimentally. Nonetheless, AI will at the very least make the work of researchers easier. As AI tools gradually assist humans in identifying new hypotheses, designing protocols, and conducting experiments, the production of relevant ideas will increase. However, the time horizon of these effects remains highly uncertain.

Once again, we draw a historical parallel to illustrate to what extent AI's ability to generate new ideas affects productivity growth. In the 17^{th} century, the invention of calculus enabled enormous advances in physics, notably in understanding the movements of projectiles or planets. Similarly, advances in glass polishing techniques allowed mankind to see the increasingly minute and thereby discover the previously unknown world of germs and other microorganisms. In the same way, AI is opening up a field of possibilities that are difficult to imagine. These effects are leading to a permanent increase in the rate of productivity growth. The magnitude of this effect, however, is difficult to quantify. When this permanent effect on productivity growth (green line in Figure 9) is added to the transitory effects of automating the production of goods and services, we obtain the blue line scenario in Figure 9.

Figure 9: Effect of AI adoption on potential productivity

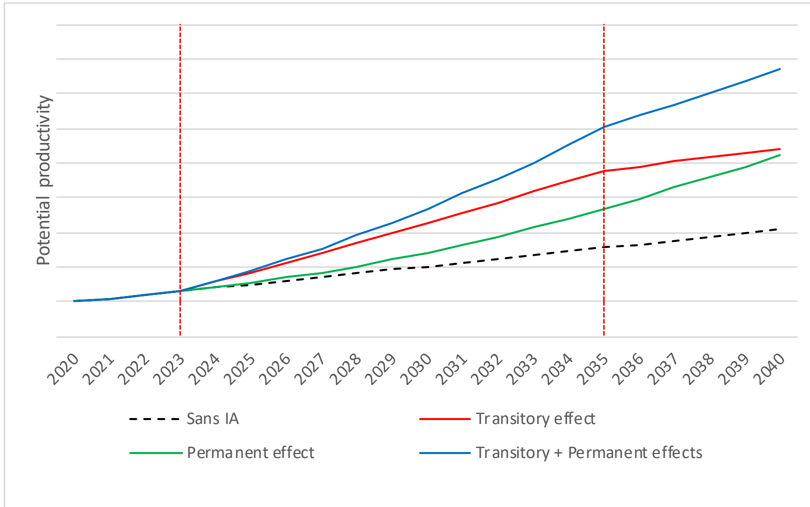

Source: French Artificial Intelligence Commission (2024)

4.2 Competition and the AI value chain

One could also argue that this estimate is overly optimistic. In the recent past, the digital revolution was supposed to lead to accelerated growth rates, but since the early 2000s, developed countries, starting with the United States, have experienced sharp declines in their growth rates. Paradoxically, at the same time, we saw the emergence of major innovations that have significantly changed our daily lives, such as the computer, the smartphone, and social networks. Several studies have offered explanations of this limited impact on growth. Some argue that it is due to a measurement issue (Byrne et al., 2016; Aghion et al., 2019), while others see it as a sign that these digital innovations have mainly improved our entertainment (Rachel, 2021). Another explanation is that the ICT revolution has fostered the emergence of superstar firms, notably the GAFAMs (Google, Amazon, Facebook, Apple, and Microsoft). While initially contributing to the observed increase in productivity growth between 1995 and 2005, overly lax competition policies allowed the GAFAMs to grow to the point where they controlled entire sectors of the economy, ultimately discouraging the entry of new, innovative firms, with negative effects on overall economic growth (Aghion et al., 2023; De Ridder, 2024). The difference between the ICT and AI revolutions is that this time the GAFAMs are dominant from the outset and can therefore immediately prevent the entry of new, innovative firms. The lack of competition is particularly pronounced in the upstream segments of the AI production chain, namely access to data and computing power, which are dominated by a small number of large firms, including the GAFAMs. It is therefore important to adapt our institutions, and in particular our competition policy, so that the AI revolution can fulfill its potential as a growth enhancer.

5 Conclusion

In this note we used two alternative approaches to estimate the macroeconomic impact of AI on productivity growth over the next decade. The first approach exploits the parallel between the AI revolution and past technological revolutions. The second approach follows Acemoglu (2024) and the task-based framework, which we revisit using our own reading of the existing empirical literature on the various component of the task-based formula. Based on the first approach, we estimated that the artificial intelligence (AI) revolution should increase aggregate productivity growth by between 0.8 and 1.3pp per year over the next decade. Using the second approach but with our own reading of the recent empirical literature on the various components of the task-based formula, we obtained a median estimate of 0.68pp additional annual TFP growth. Our estimates will be considered too pessimistic by some and too optimistic by others. The former will argue that artificial intelligence will facilitate the automation of the production of ideas, leading to a permanent increase in productivity growth (Aghion et al., 2018). The latter will point to the existence of obstacles to growth, notably the lack of competition in various segments of the AI value chain, which are already controlled by the *superstar* firms of the digital revolution (Aghion et al., 2023; De Ridder, 2024).

We see our analysis in this note as nothing more than a very first step in a broader research agenda on AI and productivity growth. A first avenue for future research will be to try and quantify the extra growth potential generated by the fact that AI automates tasks in the production of ideas. A second avenue will be to assess the importance of institutional barriers to growth, starting with the lack of competition and the inadequacy of competition policies in key segments of the AI value chain. These and other extensions of our analysis in this note are left for future research.

References

Acemoglu, Daron, "The Simple Macroeconomics of AI," *Working Paper*, 2024.

_ **and Pascual Restrepo**, "The race between man and machine: Implications of technology for growth, factor shares, and employment," *American Economic Review*, 2018, *108* (6), 1488–1542.

Aghion, Philippe, Antonin Bergeaud, Timo Boppart, Peter J Klenow, and Huiyu Li, "Missing growth from creative destruction," *American Economic Review*, 2019, *109* (8), 2795–2822.

_ , _ , _ , _ , and _ , "A theory of falling growth and rising rents," *Review of Economic Studies*, 2023, *90* (6), 2675–2702.

_ , Benjamin F Jones, and Charles I Jones, "Artificial intelligence and economic growth," in "The economics of artificial intelligence: An agenda," University of Chicago Press, 2018, pp. 237–282.

Bergeaud, Antonin, Gilbert Cette, and Rémi Lecat, "Productivity Trends in Advanced Countries between 1890 and 2012," *Review of Income and Wealth*, 2016, *62*, 420—-444.

Besiroglu, Tamay and Marius Hobbhahn, "Trends in GPU Price-Performance," Technical Report, EpochAI 2022.

Brynjolfsson, Erik, Danielle Li, and Lindsey R Raymond, "Generative AI at Work," Technical Report, National Bureau of Economic Research 2023.

Byrne, David M, John G Fernald, and Marshall B Reinsdorf, "Does the United States have a productivity slowdown or a measurement problem?," *Brookings Papers on Economic Activity*, 2016, *2016* (1), 109–182.

David, Paul A., "The dynamo and the computer : An historical perspective on the modern productivity paradox," *American Economic Review*, 1990, *80* (2), 355—361.

De Ridder, Maarten, "Market power and innovation in the intangible economy," *American Economic Review*, 2024, *114* (1), 199–251.

Dell'Aqua, Fabrizio, Edward McFowland III, Ethan R. Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Saran Rajendran, Lisa Krayer, François Candelon, and Karim R.

Lakhani, "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality," *Harvard Business School Working Paper*, 2023, (24-013).

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock, "GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models," *Technical Report, arXiv*, 2023.

French Artificial Intelligence Commission, "IA : Notre ambition pour la France," Technical Report 2024.

Gmyrek, P, J Berg, and D Bescond, "Generative AI and jobs: A global analysis of potential effects on job quantity and quality," Working Paper 96, International Labour Organization 2023.

Noy, Shakked and Whitney Zhang, "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence," *Science*, 2023, *381* (6654), 187–192.

Peng, Sida, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer, "The Impact of AI on Developer Productivity: Evidence from Github Copilot," Technical Report, arXiv 2023.

Pizzinelli, C, A Panton, M Mendes Tavares, M Cazzaniga, and L Li, "Labor Market Exposure to AI: Cross-country Differences and Distributional Implications," IMF Working Paper, International Monetary Fund 2023.

Pôle Emploi, "Les employeurs face à l'intelligence artificielle," Technical Report 2023.

Rachel, Lukasz, "Leisure-enhancing technological change," *Working paper*, 2021.

Sevilla, Jaime, Lennart Heim, Anson Ho, Tamay Besiroglu, Marius Hobbhahn, and Pablo Villalobos, "Compute trends across three eras of machine learning," in "2022 International Joint Conference on Neural Networks (IJCNN)" IEEE 2022, pp. 1–8.

Svanberg, Maja, Wensu Li, Martin Fleming, Brian Goehring, and Neil Thompson, "Beyond AI Exposure: Which Tasks are Cost-Effective to Automate with Computer Vision?," *Working Paper*, 2024.