

INTERNATIONAL MONETARY FUND

Artificial Intelligence and Productivity in Europe

Florian Misch, Ben Park, Carlo Pizzinelli and Galen Sher

WP/25/67

IMF Working Papers describe research in progress by the author(s) and are published to elicit comments and to encourage debate.

The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

**2025
APR**



WORKING PAPER

IMF Working Paper
European Department

Artificial Intelligence and Productivity in Europe
Prepared by Florian Misch*, Ben Park, Carlo Pizzinelli and Galen Sher**

Authorized for distribution by Stephan Danninger and Kristina Kostial
Month 2025

IMF Working Papers describe research in progress by the author(s) and are published to elicit comments and to encourage debate. The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

ABSTRACT: The discussion on Artificial Intelligence (AI) often centers around its impact on productivity, but macroeconomic evidence for Europe remains scarce. Using the Acemoglu (2024) approach we simulate the medium-term impact of AI adoption on total factor productivity for 31 European countries. We compile many scenarios by pooling evidence on which tasks will be automatable in the near term, using reduced-form regressions to predict AI adoption across Europe, and considering relevant regulation that restricts AI use heterogeneously across tasks, occupations and sectors. We find that the medium-term productivity gains for Europe as a whole are likely to be modest, at around 1 percent cumulatively over five years. While economically still moderate, these gains are still larger than estimates by Acemoglu (2024) for the US. They vary widely across scenarios and countries and are substantially larger in countries with higher incomes. Furthermore, we show that national and EU regulations around occupation-level requirements, AI safety, and data privacy combined could reduce Europe's productivity gains by over 30 percent if AI exposure were 50 percent lower in tasks, occupations and sectors affected by regulation.

RECOMMENDED CITATION: [Misch, F., Park, B., Pizzinelli, C. and Sher, G. (2024). Artificial Intelligence and Productivity in Europe. IMF Working Paper WP/25/67.

JEL Classification Numbers:	E24, J24, O30, O47
Keywords:	Artificial Intelligence; Productivity; Technology; Regulation
Author's E-Mail Address:	FMisch@imf.org ; HPark2@imf.org ; CPizzinelli@imf.org ; GSher@imf.org

* Corresponding author.

** We thank Simon Bunel, Era Dabla-Norris, Romain Duval and David Koll for excellent comments. Mahika Gandhi provided excellent research assistance.

WORKING PAPERS

Artificial Intelligence and Productivity in Europe

Prepared by Florian Misch¹, Ben Park, Carlo Pizzinelli and Galen Sher²

¹ Corresponding author.

² Mahika Gandhi provided excellent research assistance.

Contents

1. Introduction	3
2. Stylized Facts	6
3. Methodology.....	6
4. Results	8
4.1 Variation in Medium-term Productivity Gains	8
4.2 Preferred Scenario	12
4.3 The Role of Regulation	15
5. Conclusions and Policy Implications.....	16
Appendix 1: Sample and Country-Specific Data	18
Appendix 2: AI Exposure of Tasks and Occupations	19
Appendix 3: AI Adoption Rate	22
Appendix 4: Labor Cost Savings from AI	28
Appendix 5: National Occupation Regulation	29
Appendix 6: EU AI Act	30
Appendix 7: Data Privacy Laws	31
References.....	32

1. Introduction

Artificial intelligence (AI) is often seen as a general-purpose technology that has the potential to transform the economy and spur broad-based economic growth, akin to the arrival of electricity and personal computers. Much of the debate on its likely impact thus focuses on its effect on productivity. In Europe, this question is especially topical given lackluster productivity growth over recent decades, which resulted in a large productivity gap vis-à-vis the US (IMF, 2024). Moreover, there is a widespread view that the region is falling behind the US and China in AI development and adoption, not least because of its more stringent regulatory environment (e.g., *The Economist*, 2023). The objective of this paper is to provide estimates of the size of effects of AI on total factor productivity (TFP) across European countries over the medium term and examine any impeding effects of regulation in Europe.

Microeconomic studies suggest large productivity gains from different types of AI for specific occupations. These estimates, which in most cases are based on randomized trials where the treatment group is given access to AI tools, range from 14% for low-skilled taxi drivers to over 50% for software engineers; see Appendix 4 for a summary. Firm-level studies on the effects of adoption of AI technologies other than generative AI find smaller but still significant productivity gains (see Comunale and Manera, 2024, and Filippucci et al., 2024a for surveys). A separate strand of the literature looks at the impact of AI on employment. Cazzaniga et al. (2024), for instance, note that a large share of jobs globally is likely to be affected by AI, particularly in advanced economies, and that AI will substitute rather than complement human labor in many jobs. Using data from the US, Bonfiglioli et al. (2025) and Huang (2024) show that higher AI adoption is associated with falls in the employment-to-population ratios. A separate strand of the literature examines macroeconomic policies to broaden the gains from AI; see Brollo et al. (2024).

The extent to which these micro-level productivity gains and possible employment effects are associated with aggregate productivity gains and growth, however, remains unclear. Studies examining the medium-term macroeconomic effect of AI show a substantially wider range of estimates. McKinsey (2023) and Goldman Sachs (Hatzius et al., 2023) envision cumulative GDP gains of above 35 percent for advanced economies and 7 percent globally over a 10-year period, respectively. Commission de l'Intelligence Artificielle (2024) infers potential growth impacts from AI of up to 1.3 annually by drawing parallels to the effects of electricity and Information and Communication technologies. Similarly, IMF (2024) and Cazzaniga et al. (2024) estimate annual growth impacts of up to 0.8 percentage points based on labor reallocation and changes in the capital share comparable to those observed for automation in the past.

By contrast, Acemoglu (2024) does not take into account any potential longer-term transformational effects of AI and therefore estimates much more modest TFP gains of less than 0.7 percent cumulatively over 10 years which he refers to as 'medium term'. He uses a rigorous framework that quantifies the gains bottom-up using measures of the AI exposure of individual tasks for each occupation. Aghion and Bunel (2024) show that using alternative assumptions within Acemoglu's (2024) framework can 10-fold the estimated productivity gains for the US.

Evidence considering Europe-specific factors and cross-country heterogeneity within the region remains relatively scarce. Bergeaud (2024) simulates productivity gains from AI for the euro area using the Acemoglu (2024) framework, combining some of the original paper's parameter values with own assumptions and estimates to generate a few alternative scenarios. He finds cumulative productivity gains of 2.9 percent for the

euro area in his central scenario, while his country-specific results range from around 1.5 in Ireland to 3.3 in Belgium. Filippucci et al. (2024b) extend the Acemoglu (2024) framework by more explicitly modelling sectoral spillovers and present the productivity gains for all G7 countries while assuming different adoption rates. The estimated gains are also significantly higher than the Acemoglu (2024) results for most countries and scenarios.

Our study is broader in scope in terms of country coverage, number of scenarios and consideration of the effects of regulation, and it uses econometrically-grounded parameter estimates. We investigate cross-country variation for 31 European countries both in terms of the magnitude and the uncertainty of the productivity gains more systematically by allowing the rates of AI adoption to vary across countries according to their economic characteristics. To this end, we also use Acemoglu's (2024) framework to estimate the medium-term productivity gains from AI in Europe (which we interpret to be 5 years, given the model characteristics and our assumptions).¹

First, we quantify the uncertainty around the impact of AI on TFP. Rather than making our own assumptions, we combine a comprehensive set of the available estimates of AI exposure of individual tasks, delivering 44 scenarios. For each scenario, we calibrate estimates of AI adoption by Svanberg et al. (2024), as used in Acemoglu (2024), to specific countries and sectors based on regression evidence of the drivers of AI adoption in Europe. Wage levels turn out to be the main driver of AI adoption, rather than capital costs, industry concentration, digitalization, or human capital. This allows us to take into account how differences in wage levels and other sectoral and country-level factors affect AI adoption and shape cross-country variation of the productivity gains.

Second, we examine the role of regulation in reducing productivity gains of AI which is one area of policy that is often discussed when it comes to AI; see for example Bradford (2024) for a broader discussion. Survey evidence from Germany suggests that while firms see many barriers to AI adoption, regulation is seen as most important (Wintergerst, 2024). To this end, we consider regulation which could plausibly have large effects on AI adoption: licensing and training requirements for specific occupations at the national level, data privacy laws, and the EU AI Act.² We then identify the tasks, occupations and sectors where this regulation could undermine AI adoption and assume that regulation halves AI capabilities, striking a reasonable middle ground between assuming that regulation completely prevents AI use and that regulation has no impact.

We show that in our preferred scenario, which is based on what we think are the most plausible assumptions on AI occupational exposure and adoption rates in European countries, the Europe-wide effects are modest at around 1.1 percent cumulatively over the medium term, which exceeds Acemoglu's estimates for the US by almost 60 percent. These differences are driven almost entirely by more optimistic assumptions of AI capabilities which we think are justified in light of a recent refereed publication which we explain below. However, there is significant heterogeneity across countries. Estimated TFP gains in higher-income countries tend to be much larger than those in lower-income economies in line with findings by Cerutti et al. (2025), due

¹ In contrast to Acemoglu (2024) who considers a 10-year period, we assume that the simulated productivity gains refer to a 5-year horizon for two reasons. First, the framework by Acemoglu (2024) does not capture any large-scale transformational effects which we think could indeed arise over periods exceeding 5 years. Second, we assume that AI capabilities and the AI adoption rate (i.e., the proportion of tasks for which it will be profitable to use AI) will remain constant over time. Over 10 years, both are more likely to change.

² The EU AI Act is a key AI safety law which caps the capacity of AI systems and increases the cost of using AI in a defined list of high-risk systems.

to both their larger share of value added in industries that have greater AI exposure (e.g., financial services) and their higher wage levels—which provide greater incentives for labor-saving or labor-augmenting AI adoption.

Moreover, in higher income-countries, there are much larger upside risks from AI. For example, in our preferred scenario, the gains in Luxembourg could be 2 percent cumulatively, almost twice the European average, and more than 4 times larger than those in Romania. This is due to the composition of the Luxembourgish economy, with more value added in sectors like financial services with higher AI exposure, and due to Luxembourg's higher wages, which give employers there a greater incentive to adopt AI. In addition, productivity gains in Luxembourg could be more than twice as high if AI turns out to be more capable than in the 'preferred' scenario, pointing to larger upside risks as well.


We also find that the combined adverse effects of national occupation-level regulation, the EU AI Act and data privacy laws on productivity could be significant, with the former two having the largest effects. Occupational restrictions are often overlooked in debates on AI but could substantially reduce the productivity gains from AI. By contrast, data privacy laws widely affect some industries with high AI exposure, such as the IT and financial services sectors, but their productivity-inhibiting effect is somewhat smaller.

Our results can inform ongoing policy debates. On the one hand, they suggest that over the medium term, AI is not a silver bullet to significantly boost sluggish productivity growth in Europe, or to close the productivity gap to the United States. Moreover, AI could slow the rate of income convergence among European countries because the gains from AI tend to be larger in more advanced economies. On the other hand, our results help inform the potential trade-offs between the benefits of regulations including related to privacy and safety, and maximizing the productivity gains from AI, while noting that our study is not a comprehensive cost-benefit analysis of these regulations.

This paper is organized as follows. Section 2 presents stylized facts on the diffusion and adoption of AI to motivate the analysis and the focus on medium-term effects. Section 3 provides a summary of the Acemoglu (2024) model and discusses how we calibrate some of its key parameters to the European context, considering a range of alternative scenarios. Section 4 presents the results and the analysis of regulations. Section 5 concludes.

2. Stylized Facts

The fast speed of diffusion of generative AI (genAI), a subset of the broad set of technologies that fall under the definition of AI, over the last two years explains in part the recent interest in and public debates on the effects of AI more broadly. Compared to past innovations, it took only a few months for genAI to reach 100 million users (measured by the user base of OpenAI's ChatGPT which was the first widely available and accessible genAI application). By contrast, it took years (and sometimes decades) for other general-purpose technologies to reach the same number of users. Even though these differences may be mainly driven by the very low access cost of genAI technologies through smartphones and personal computers, this historically unprecedented speed of adoption points to the potential for AI to be applicable to a wide range of tasks. However, as we discuss below, the sheer number of individual users neither implies that AI is being employed to a wide range of tasks nor is it synonymous for broader AI adoption by firms.



3. Methodology

To simulate the effects of AI on productivity, this paper uses the model from Acemoglu (2024), which in turn is based on Acemoglu and Restrepo (2018, 2019, 2022). In the model, the production of a unique final good requires a fixed set of tasks to be performed, and in turn, these tasks can be produced with either capital or labor. This framework serves well to examine the medium-term effects of AI, interpreted as the growth in output and TFP originating from small changes in productivity and in the mix of labor and capital inputs, but without fundamental long-run changes in the structure of the economy (i.e., its sectoral composition and tasks).

Acemoglu (2024) includes two channels for AI-based productivity gains: automation and task complementarity. The former entails a substitution of workers in the performance of individual tasks within an occupation, decreasing the overall need for human labor. The latter refers to partial automation of tasks so that the marginal labor productivity increases in complementary tasks performed by humans.

Acemoglu (2024) then applies Hulten's theorem (Hulten, 1978), which shows how micro-level, non-transformational productivity improvements in individual occupations translate into macro changes and aggregated productivity growth in competitive economies with constant returns to scale. Since Acemoglu's focus is on small changes in technology and the competitive equilibrium is efficient, the impact of all reallocations of factors across tasks and indirect effects via prices are of secondary order and therefore small enough to be ignored in computing the productivity and GDP gains due to AI. This also implies that the model

is not suitable for estimating any potentially large transformational and longer-term effects, such as the creation of new industries or an acceleration in the rate of scientific discoveries.

Estimating the productivity gains of AI adoption in the Acemoglu (2024) model amounts to calibrating three key parameters. The first required input consists of a measure of the exposure to AI of different occupations, essentially reflecting assumptions about AI capabilities and the potential scope of AI applications. Importantly, for the purpose of estimating productivity gains, exposure can refer to both automation and task complementarities.³ Using a version of the task-based occupational exposure measure constructed by Eloundou et al. (2024), Acemoglu (2024) calculates that the wage bill-weighted share of exposed tasks in the U.S. economy is 19.9 percent, meaning that AI has the capability of performing around 20 percent of tasks in the U.S. economy (when wage bill weights are used to proxy the relative importance of tasks and occupations).

The second parameter of interest is the AI adoption rate, i.e., the share of AI-exposed tasks where the benefits of using AI exceed the costs, thus making AI adoption profitable. Intuitively, there may be tasks that AI can perform but where its application may be too costly relative to the price of labor in the same job. Hence, the raw exposure measure of an occupation should be combined with an estimate of the economic feasibility of adoption. Acemoglu (2024) draws from the costing estimates of Svanberg et al. (2024) to assume that the benefits of using AI exceed the costs for 23 percent of AI-exposed tasks.

Finally, translating the application of AI into productivity gains requires an estimate of the savings in terms of labor costs that AI provides when producing a unit of output. Acemoglu (2024) uses an average of three microeconomic productivity estimates to calibrate the assumption that tasks automated by AI reduce labor costs by 27%. These labor cost savings translate into 15% total (labor and capital) cost savings, given that labor costs account for 53 percent of output. Multiplying these three key parameters (19.9%, 23%, and 15%), Acemoglu (2024) estimates a 0.71% cumulative medium-term increase in total factor productivity for the US. In contrast to Acemoglu (2024) who considers the medium term to be 10 years, we assume that the cumulative medium-term gains arise over 5 years.

We apply the same methodology to 31 European countries (see Appendix 1 for list of countries). For each country, we obtain the first key parameter—the share of AI-exposed tasks in industry-specific value added—by weighting occupation-level exposures by each occupation's wage bill within an industry (see Appendix 2 for the AI exposure estimates from the literature we use). To obtain the second key parameter—the AI adoption rate—we estimate the historical relationship between AI adoption and countries' and sectors' economic characteristics (e.g., labor and capital costs) in Europe. This gives the share of tasks in each country-industry pair for which it is profitable to apply AI (see Appendix 3). We calibrate the third key parameter (cost savings) in the same way as Acemoglu (2024). Finally, we calculate productivity gain as the product of these three key parameters (see Appendix 1 for details on data used for weighting and labor shares).

³ The distinction between automation and complementarity, however, would be relevant for examining the impact of adoption on employment and labor productivity, which we leave for future research.


4. Results

4.1 Variation in Medium-term Productivity Gains

We present the results in several stages. First, we compare productivity gains in Europe to the estimates by Acemoglu (2024) for the US. Second, we show how alternative assumptions of AI capabilities and AI adoption affect the productivity gains across the countries in our sample. Finally, we combine all assumptions to arrive to a plausible range of estimates for each country.

First, we apply the same methodology as in Acemoglu (2024) to Europe in order to compare our estimates with the ones for Acemoglu (2024) for the US. In particular, we calculate the productivity gains for European countries and Europe as a whole, using the same measure of occupation-level AI exposure, AI adoption rate, and labor cost savings from AI as in the original Acemoglu (2024) paper. This parametrization, henceforth referred to as the ‘Acemoglu (2024) baseline’, allows us to examine how differences in the sectoral composition and wage structure translate into differences in the productivity gains from AI between the US and European countries, given that there are no other differences between Acemoglu’s result and our results for Europe.

Figure 2 shows that, based on differences in the sectoral composition, the average productivity gains in Europe are somewhat larger than in the US, but not by much: they amount to close to 0.8 percent in Europe, cumulatively over the medium term, compared to around 0.7 in Acemoglu (2024) for the US. However, there is substantial cross-country variation, ranging from around 0.5 percent in Romania to close to 1 percent in Luxembourg. Broadly, higher-income countries have larger gains, a result driven by the higher prevalence of white-collar services including for instance financial services, which tend to be more exposed to AI.


Figure 2. Productivity gains in Europe and the US compared

Sources: Eurostat, EU LFS, EU SILC, and IMF staff calculations.

Note: The black and red lines represent the average TFP gain for the 31 European countries in our sample and for the US as estimated by Acemoglu (2024), respectively.

Second, we quantify the uncertainty around the Acemoglu (2024) baseline results presented in Figure 2 using a comprehensive set of alternative scenarios about AI capabilities (AI task-level exposure measures).⁴ We thus repeat the exercise considering various alternative estimates on AI exposure. Appendix 2 provides a list of AI exposure estimates and brief descriptions of their main methodological differences. For instance, some compute exposure for specific tasks and then consider each occupation as a bundle of tasks (Eloundou et al., 2024, Gmyrek et al., 2023, Webb, 2019), while others focus on the overlap between AI application and human skills (Felten et al., 2021). Beyond the mechanical meaning of testing the sensitivity of the results to these assumptions, exploring alternative measures of exposure thus also reflects the considerable uncertainty around how individual tasks and occupations will be affected by AI. Figure 3 shows that while more conservative assumptions mute the productivity gains of AI and its variation across countries, more optimistic assumptions suggest that there are large upside risks to productivity, especially in countries that are set to gain more in the Acemoglu (2024) baseline. For instance, in Luxembourg, alternative assumptions imply that productivity gains could be above 3 percent, cumulatively over the medium term, more than 2 percentage points higher than under the Acemoglu (2024) baseline.

⁴ Some of the AI exposure measures only refer to generative AI. For simplicity, we assume that all of the exposure measures broadly refer to AI technologies in general.

Figure 3. Productivity gains from AI under different assumptions about AI exposure at the task level

Sources: Eurostat, EU LFS, EU SILC, and IMF staff calculations.

Note: Each set of dots represents the estimated cumulative medium-term TFP impacts from AI adoption based on alternative measures of occupation-level AI exposure. 'Europe' refers to the average of the countries in our sample.

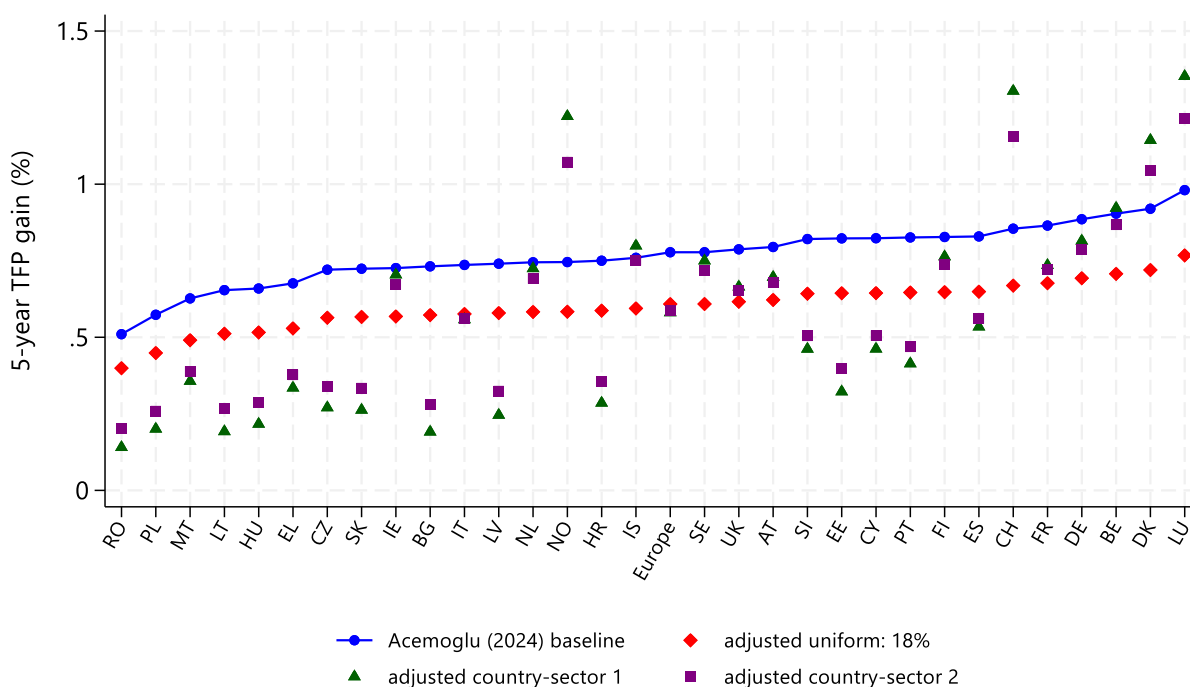
Third, we use the same assumptions for AI task exposure and labor cost savings as in the Acemoglu (2024) baseline but vary the estimates of the rate of AI adoption. This is a data-driven extension to the Acemoglu (2024) baseline, which assumes that AI is cost-effective to adopt for 23% of firms in every country and sector based on Svanberg et al. (2024). Using country-sector-level regressions that examine possible drivers of AI adoption, we re-calibrate Svanberg et al.'s (2024) estimate of the rate of AI adoption based on the evidence for how AI adoption in Europe has varied historically across countries and sectors, as explained in Appendix 3. This differs from the approach taken by Filippucci et al. (2024b) who use variation in the AI preparedness index compiled by Cazzaniga et al. (2024) to infer variation in the country-specific adoption rates. In Appendix 3, we also discuss the possibility and implications of much lower costs of AI adoption.

This analysis suggests that wages are a robust factor explaining AI adoption rates historically across European countries and sectors, suggesting that firms in high-wage countries and sectors have more incentive to adopt AI, either to automate certain tasks previously done by workers or to boost their workers' productivity at performing complementary tasks. Given that Europe has lower wages than the US, these estimates suggest that AI adoption rates could be around 5 percentage points lower in Europe than in the US. Therefore, we calibrate one adoption rate scenario where AI adoption is 18 percent in Europe, 5 percentage points lower than under the Acemoglu (2024) baseline for the US. This scenario assumes that the adoption rate is the same across all European countries and sectors. Then, we consider two alternative scenarios that allow the rate of AI adoption to vary across European countries and sectors according to their wage rate, both of which maintain the average Europe-wide adoption rate of 18 percent. The two scenarios differ only on a technical level in how

they ensure that the average Europe-wide adoption rate is maintained at 18 percent – one scenario shifts the wage-based estimates of AI adoption rates upward from their 2023 levels and the other rescales those estimates upward (see Appendix 3).

Figure 4 shows the productivity gains under alternative estimates of AI adoption. Overall, when AI adoption rates are allowed to vary across countries and sectors, the range of estimated productivity gains across European countries increases. This is because high-wage countries, which have stronger incentives to adopt AI, are also those countries where more of the tasks can be performed by AI, because their economic activity is composed more of sectors like services that are highly exposed to AI. While the effect of using different measures of exposure was monotonic (i.e., the productivity gains would fall or increase for all countries), calibrating country-sector-specific AI adoption rates yields greater productivity gains for some (mainly higher-income) countries and lower gains for others (mainly lower-income countries). Intuitively, countries seeing upward revisions in productivity gains are those where wages are higher (e.g., Luxembourg, Norway, and Switzerland).

Figure 4. Productivity gains from AI under different assumptions about the AI adoption rate




Sources: Eurostat, EU LFS, EU SILC, and IMF staff calculations.

Note: Each set of dots represents the estimated cumulative medium-term productivity gains from AI based on alternative measures of AI adoption rates across sectors and countries, as explained in Appendix 3. 'Europe' refers to the average of the countries in our sample.

Finally, we combine all assumptions with respect to occupational AI exposure and with respect to AI adoption to arrive at 44 different scenarios. This exercise quantifies the uncertainty in the estimated productivity gains for Europe around the Acemoglu (2024) baseline and shows the sensitivity of these estimates to the assumptions. Figure 5 shows that productivity gains vary significantly both between European countries and within countries across the different scenarios. Interestingly, the cumulative medium-term productivity gains are more uncertain

in countries that are expected to gain more. For example, in Norway, which is expected to gain the most in the median scenario, there is one scenario where its productivity gain amounts to around 5.1% cumulatively over the medium term, while more conservative scenario assumptions place the effects near zero. Luxembourg therefore has the largest inter-quartile range across scenarios, spanning from 0.8% to 2.95%. Meanwhile, Romania, the country with the lowest productivity gain in the median scenario, has the lowest uncertainty around this estimate—its interquartile range is from approximately 0.3% to 0.7%.

Figure 5. Overview of the productivity gains from AI in all scenarios



Sources: Eurostat, EU LFS, EU SILC, and IMF staff calculations.


Note: The box plot reports, for each country, the distribution of the estimated cumulative medium-term productivity gains across all combinations of AI exposure measures and adoption rates. The middle line of the box reports the median productivity gain, while the lower and upper borders of the box report the 25th and 75th percentiles of the distribution of productivity gains, respectively. The upper (lower) whiskers report the upper (lower) adjacent values, that is the largest (smallest) value lying below (above) 1.5 times the inter-quartile range from the 75th (25th) percentile. The 'preferred' scenario is explained in Section 4.2. 'Europe' refers to the average of the countries in our sample.

4.2 Preferred Scenario

While Figure 4 captures the large degree of uncertainty in the estimated productivity gains, out of these 44 scenarios, we select our preferred scenario as the combination of AI occupational exposure measure and adoption rate that we think best aligns with plausible assumptions and characteristics of European countries. With respect to the measure of AI exposure of occupations, we select the baseline task-based estimates from Eloundou et al. (2024), a widely cited measure published in a top-tiered journal. We use the scenario of AI

adoption rates that vary across countries and sectors according to wages, while shifting the estimated sector-country adoption rates upward to keep the average adoption rate in Europe at 18%; this scenario is referred to as ‘adjusted country-sector 1’ in Figure 4.


The diamonds in Figure 5 show the productivity gains under the preferred scenario which vary significantly across countries, due to differences in the countries’ economic structure and the wage-driven AI adoption rates. Figure 6 shows that on average, the medium-term productivity gains from AI in Europe based on the preferred scenario are around 1.1 percent, almost 60 percent larger than those under the Acemoglu (2024) baseline.



Finally, in Figure 7, we decompose the productivity gains by showing what drives the deviation for each country from the European average of the productivity gains under the preferred scenario. Within our framework and the assumption of homogenous occupation-level AI exposure across countries, there are three possible drivers of cross-country differences: the AI adoption rate (which varies across countries and sectors in the preferred scenario), the industry composition of value added, and the occupational composition within each industry. Figure 4 shows that heterogeneity in adoption rates, due to the relative cost of labor across countries and sectors, drive most of the cross-country variation in our estimated productivity gains in the preferred scenario.

Differences in countries' industrial composition is the second-largest driver. In particular, countries with the highest expected productivity impact are characterized by both higher adoption rates and higher concentration of value added in high-exposure industries. Finally, cross-country differences in occupational composition of individual industries account for only minimal variation across countries.

Given non-linearities, the framework is not additive, and the three drivers that we single out in Figure 7 do not equal the exact value of the preferred scenario. Hence, the gray bars in Figure 7 report the residual component—that is, the difference between the sum of the three drivers (industry composition, occupation composition, and AI adoption) and the gains in the preferred scenario. In most cases, the residual is negative, especially for countries with the largest estimated productivity gains, implying that the sum of the three drivers overpredicts the gains in the preferred scenario.



4.3 The Role of Regulation

In this subsection, we turn to the role of regulation in shaping the productivity gains of AI. We focus on three types of regulation:

Regulated occupations. There are a host of licensing and training requirements for specific occupations implemented by individual countries (referred to as national occupation-level regulation). We use the European Commission Regulated Professions Database to determine which occupations in which countries are subject to this type of regulation. Appendix 5 provides details. While this regulation pre-dates the recent spread of AI technologies, we still assume that it essentially discourages the use of AI for these occupations, not least because trained and licensed workers are required to perform certain occupations by law.

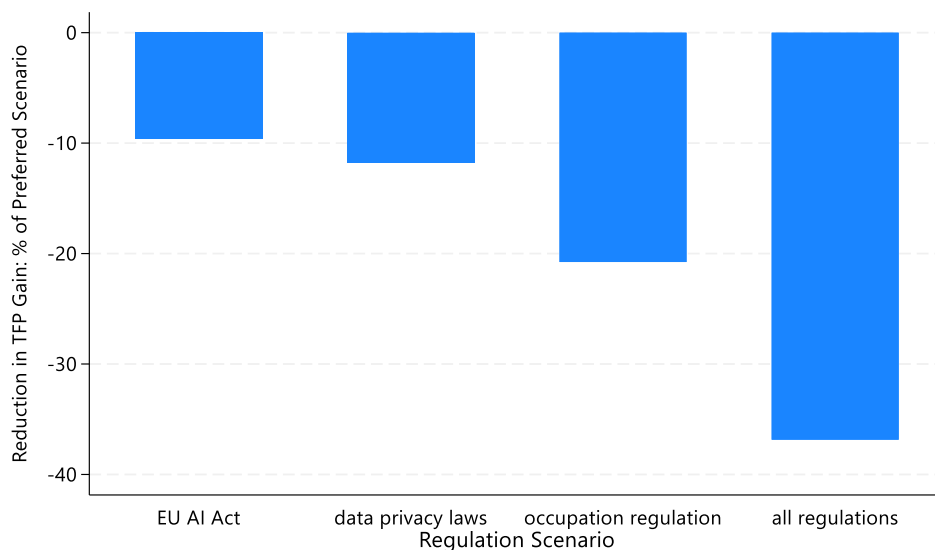
EU AI Act. Introduced in 2024 to enhance AI safety, the EU AI Act could affect the productivity gains of AI in two ways. First, the EU AI Act specifies regulatory requirements for general purpose AI models that exceed a defined threshold of computational power. The Act views these models as systemically risky because of their far-reaching applications. However, the threshold set by the now revoked US *Executive Order on the Safe, Secure, and Trustworthy Development and Use of AI* was some magnitudes higher, implying that the threshold imposed by the EU AI Act governs a broader set of models (Larsen and Küspert, 2024). European Commission (2024) points out that some existing and publicly accessible models already exceed that level. To model the impact of the EU AI Act, we draw on Acemoglu's (2024)'s extension of his simulations that differentiates between tasks like coding or summarization that are 'easy' for AI to learn and perform, and those that are 'hard', like diagnosing and treating a patient in a medical setting. Rather than assuming that AI systems above the threshold defined by the EU AI act matter for all tasks, we only assume that they might be used for tasks that Acemoglu specifies as 'hard', and that the additional reporting and oversight requirements would discourage the use of AI for these tasks classified as 'hard'. The second way that the EU AI Act could affect the productivity gains of AI is in the regulatory requirements that it places on tasks that it views as 'high-risk'. Here, we likewise assume that the additional regulatory burden discourages the use of AI in occupations that perform such tasks. In Appendix 6, we provide details of how these assumptions are introduced in the framework.

Data privacy laws. We consider the effects of data privacy laws more broadly without referring to specific European or national regulation. We assume that data privacy regulation complicates and therefore discourages the use of AI in three data-intensive sectors: Information and Communication (NACE Rev. 2 code J), Financial and Insurance Activities (NACE Rev. 2 code K), and Human Health and Social Work Activities (NACE Rev. 2 code Q).

While we do not know the exact extent to which occupation regulation, the EU AI Act and data privacy laws plausibly undermine AI use in regulated occupations, for hard and 'high risks' tasks, and in data-intensive sectors. For illustrative purposes we compute the effects if the AI exposure will be half of what they would be otherwise; i.e., we lower AI exposure by 50 percent in those tasks, occupations and sectors affected by regulation. In our view, this strikes a reasonable balance between assuming that AI continues to have a full impact despite regulation and assuming that regulation completely stifles AI use. Figure 8 reports the result by type of regulation and the combined effect of all three regulatory areas relative to the preferred scenario. It shows that data privacy laws reduce the productivity gains of AI by 10 percent, while the EU AI Act and national

occupation-level regulation each lower the productivity gains by around 15 percent. The combined effect of these regulations is substantial, reducing the productivity gains from AI by well over 30 percent. Note that these scenarios are hypothetical as not all of these regulations are fully implemented and applicable in all countries of our sample.

Figure 8. Productivity Gains from AI under different assumptions about the AI exposure at the task level



Sources: Eurostat, EU LFS, EU SILC, and IMF staff calculations.

Note: The first three bars report the fall in the estimated short-run TFP impact under each policy scenario compared to the EUR average. The final bar reports the respective fall under a combined scenario in which all policies apply.

5. Conclusions and Policy Implications

This paper estimates the cumulative productivity gains from AI for European economies over the medium term. The estimates are based on the Acemoglu (2024) model, which is a transparent and intuitive framework designed to capture incremental productivity changes, rather than larger longer-term transformational effects of AI adoption, such as a faster pace of scientific progress. Relative to the existing literature, this paper quantifies uncertainty by compiling different scenarios, examines in detail cross-country variation taking into account differences in wages and other factors that could drive AI adoption rates at the country-sector level, and estimates the effects of regulation on the productivity gains from AI if AI exposure were 50 percent lower in tasks, occupations and sectors affected by regulation.

We find that the productivity gains from AI in Europe are likely to be modest, at around 1.1 percent over the medium term in our preferred scenario. This estimate is still small but almost 60 percent above the estimates for the US from Acemoglu (2024). The main reasons are differences in Europe's sectoral composition and different assumptions about AI capabilities. We also show that there is significant variation in the productivity gains from AI across European countries. Higher-income countries are predicted to benefit more, both because they have larger sectors like white-collar services that benefit most from AI and because their higher wages create stronger incentives for their firms to adopt AI. Therefore, AI could slow down the rate of income

convergence within Europe. This macroeconomic result could be seen as the opposite of the finding, at the microeconomic level, that AI boosts the productivity of low-skill workers (which are probably more prevalent in lower-income countries) far more than that of higher-skilled workers (Brynjolfsson et al., 2023). Although higher-income countries look set to enjoy higher productivity gains from AI, they also face larger uncertainty around these potential gains.

Our results are subject to various caveats. The costs of AI systems are often seen as very high, given the large costs of chips and energy needs (Nathan et al., 2024). However, if the costs of deploying AI systems fall more quickly or are lower than estimated in the Acemoglu (2024) baseline, for example if many business-to-business AI service providers are created, the rates of AI adoption and productivity gains could be much larger. In this case, Svanberg et al. (2024) estimate that the rate of AI adoption in the US could be around 80%, around three-and-a-half times as large as the 23% assumed in the Acemoglu (2024) baseline; the associated productivity gains would then also increase to the same extent, as discussed in greater detail in Appendix 3. Along the same lines, the labor cost savings could be larger than what Acemoglu (2024) and we assume if alternative estimates were chosen from the range of available microeconomic estimates, which we review in Appendix 4. Finally, it could take longer for our estimated productivity gains to materialize, not least because some of the exposure estimates assume some degree of improvements relative to existing AI technologies.

It is also important to stress over the longer term the productivity gains are likely to be larger than our estimates due to the potential for AI to lead to structural transformation of the economy. AI could create new industries and value chains and accelerate R&D activity through generating novel research ideas and investigating them (Si et al., 2024) and local spillovers from the development of AI itself, all of which we leave to future research. Considering these transmission channels of AI on productivity could lead to more permanent growth effects, rather than one-off level effects from automating existing tasks (Frey et al., 2024).

Appendix 1: Sample and Country-Specific Data

This Appendix describes the data sources and other details of data preparation for the simulations. We apply the framework of Acemoglu (2024) to data from 31 European countries, including the EU27 countries, Iceland, Norway, Switzerland and the UK. To aggregate task-level AI exposures to the industry level, we compute weights for all occupations within each industry and country, using country-industry-occupation level hours worked and country-occupation-level wages. We then use industry value added and the labor share of each industry to arrive at macroeconomy-wide estimates of the share of AI exposed value added which could be subject to labor cost savings. The average of our sample of AI exposed value added in our preferred scenario is 6.8%.

Hours at the industry-occupation level

For most countries, we use data from the EU Labour Force Survey (EU-LFS) to obtain hours worked by occupation within each industry. Workers' occupation of employment is generally recorded at the ISCO-08 3-digit level, with the exception of Malta and Luxembourg, where occupations are reported at the 1-digit level. For the Netherlands and the UK, we use publicly available data at the 1-digit ISCO-08 and NACE Rev. 2 level provided by Eurostat.⁵ We include all sectors except agriculture, activities of households and employers, and activities of extraterritorial organizations. The data refers to year 2023 for all countries except for the UK, for which we use 2019, the latest available on Eurostat's website.

Wages at the occupation level

We use the EU Survey on Income and Living Condition (EU-SILC) to obtain wage data by occupation for most countries in the sample. For the majority of these countries, occupations are reported at the 2-digit level of the ISCO-08 classification. For a few countries, only 1-digit level occupations are available. Average hourly wages by occupations are computed by dividing yearly labor income by the number of weekly hours worked multiplied by 52 weeks. Although industry information is available, due to sample size limitations, we do not compute average hourly wages by occupation and industry. For the Netherlands and the UK, we use publicly available tabulations produced by Eurostat based on the Structure of Earnings Survey, referring to 2019.

Wage bill at the occupation-industry level

To compute the wage bill of each occupation in an industry, we multiply the occupation-specific wages with the occupation-industry-specific hours worked. The implicit simplification is that occupations which exist in several industries are remunerated identically, i.e., they have the same wage in a given country. Moreover, as discussed above, the wage data from the EU-SILC is at a more aggregate level than the hours data from the EU-LFS. Therefore, multiple 3-digit occupations from the EU-LFS are merged with their higher-level corresponding 2-digit code from the EU-SILC (which is sufficiently granular to capture any substantial variation across occupations). The granularity of the data remains unchanged as AI exposure estimates are available at the 3-digit occupation level.

National accounts: value added and labor share

We use national accounts statistics, publicly available on Eurostat's website, to obtain value added and labor share of compensation by industry at the 1-digit NACE 2 level. For all countries except the UK, data refers to

⁵ Access to microdata for the Netherlands was not granted under the application submitted to Eurostat for this project. The UK data was discontinued from Eurostat after the country's exit from the EU.

2023, while for the UK, data refers to 2019. Country-level labor shares of compensations are computed as averages of the sectoral shares weighted by the sectors' value added. These are needed to compute the saving from AI adoption as a share of GDP, as the reduction in production costs from using AI only concerns labor costs. For instance, for the US, Acemoglu (2024) multiplies the estimated 27 percent cost saving from AI adoption from reduced labor input by the labor compensation's share of US GDP of 0.535.⁶

Appendix 2: AI Exposure of Tasks and Occupations

For the baseline result and the robustness checks, we apply several alternative estimates of task- and occupation-level exposure to AI. In this Appendix, we include a concise description of the various estimates, their key differences, and how we adjust them to use them in our exercise. In some cases, we convert the measures from the occupational classifications in which they are originally available (e.g., the US SOC 2010 or the O*NET SOC 2018) to ISCO-08 using crosswalks published by the Bureau of Labor Statistics. Many-to-one matches are averaged using simple means.

Each source considered in the analysis takes a different approach to measure exposure and, as discussed in the main text, results in more or less conservative estimates of the scope for AI adoption. While this is to some extent driven by the approaches themselves, it may also result from choices of cutoffs used to define occupations and tasks as exposed or not exposed. The measures considered are listed in Appendix Table A1. Details on the definition of exposure of each measure and how we adapt it to our analysis are as follows:

- **Eloundou et al. (2024):** Exposure is measured at the task-level. In general terms, exposure is interpreted as whether performing of a task within a given occupation would take at least 50 percent less time if performed using AI. Task-level values are then averaged across all tasks within an occupation to obtain the occupation's share of tasks that is exposure to AI. The authors present alternative measures of exposure, either scored by humans or by a generative AI model and frame the questions in terms of either currently available AI technologies or some degree of additional technologies yet to be invented. As a default measure, we use their machine-scored β measure, which assume some degree of additional innovation. For the "low" version we use their machine-scored α measure, which assumes no additional innovation and is thus more conservative on the potential application of AI. For the "high" version, we use the machine-scored γ values, which fully assume the presence of additional innovation.
- **Acemoglu (2024):** Exposure is measured at the task level. The author constructs this measure starting from an alternative version of the Eloundou et al. (2024) index, where the β exposure is measured on a scale of 0, 0.25, 0.5, 0.75, or 1. Acemoglu (2025) recodes this measures and assigns 0 if the original value is equal to or less than 0.5 and assigns 1 if the original score is 0.75 or 1. In addition, Acemoglu (2025) also provides an alternative exposure measure where AI-exposed tasks are subsequently divided into 'easy' and 'hard' ones. In a robustness check, Acemoglu (2024) assumes lower cost savings for 'hard' tasks (7 instead of 27 percent for 'easy' tasks).

⁶ See footnote 24 on page 28 of Acemoglu (2024).

- **Felten et al. (2021):** Exposure is measured as the degree of overlap between a set of AI applications (e.g., image recognition, text creation) and the essential skills needed to perform a job, weighted by their importance and relevance scores as reported in O*NET. The latter is a large occupation data repository for US jobs which contains essential skills for each occupation. Since the original score is a continuous value, we turn it into a binary variable at the occupational level by choosing a threshold value to define exposed occupations. Since the choice of such value is somewhat arbitrary but fundamentally represents a given degree of optimism for the applicability of AI, we consider two alternative thresholds (both of which are plotted in Figure 3): (i) its median value across all ISCO-08 occupations (reported as ‘Felten et al (2021) – 50th pct.’) and (ii) the 75th percentile (reported as ‘Felten et al (2021) – 75th pct.’). Intuitively, using the 75th percentile represents a more conservative estimate of the scope for AI to replace human labor in performing essential tasks of a job, and thus implies a smaller share of exposed occupations and consequently a lower potential increase in TFP (see Figure 3).
- **Gmyrek et al. (2023):** Exposure is measured at the occupation level as the share of tasks that can be automated by AI, with individual continuous task-level scores assigned by a generative AI model. Occupations are then assigned to different categories based on their distributions of task-level exposure score: ‘Not Affected:’ AI is unlikely to have tangible effects; ‘Automation Potential’ AI is likely to substitute labor; ‘Augmentation Potential’ AI is likely to complement labor; and ‘The Big Unknown’ it is hard to predict whether AI will have an automating or augmenting effect or no effect at all. Occupations are assigned to these categories based on their distributions of task-level exposure score. We use the raw average task-level score within an occupation to represent exposure. In our alternative “strict” measure, we recode occupation-level exposure as 0 if the occupation belongs to the “No Impact” or “The Big Unknown” categories.
- **Webb (2019):** Exposure is measured through an index representing the overlap in salient word-noun pairs contained in the text description of occupation’s tasks, as reported in O*NET and the description of AI-related historical patent filings. This measure has the advantage of restricting exposure to the contexts to which AI is most plausibly being applied to, as proxied by the filing of a patent, rather than the potential scope of the technology. However, using data on patents filed makes the measure more backward-looking as it cannot capture future capabilities of AI. Similar to our approach with the Felten et al. (2021) measure, we consider both the 50th and the 75th percentiles of the raw exposure score across occupations as a cutoff to define AI-exposed occupations. The latter is particularly restrictive and, as seen in Figure 3, it effectively leads to no exposure in countries for which the underlying occupational data is at the 1-digit ISCO-08 level, as the vast majority of major occupation groups have a score below the 75th percentile.

Appendix Table A1: Alternative AI exposure measures

Measure	Source	Details
Acemoglu (2024) Baseline	Acemoglu (2024)	Baseline measure from the original paper
Acemoglu (2024) - Easy-Hard Tasks	Acemoglu (2024)	Easy-Hard task extension from the original paper
Eloundou et al. (2024)	Eloundou et al. (2024)	GPT-estimated β measure
Eloundou et al. (2024) - Low	Eloundou et al. (2024)	GPT-estimated α measure
Eloundou et al. (2024) - High	Eloundou et al. (2024)	GPT-estimated γ measure
Felten et al. (2021) - 50th pct.	Felten et al. (2021)	Exposed occupation: AIOE > 50th percentile
Felten et al. (2021) - 75th pct.	Felten et al. (2021)	Exposed occupation: AIOE > 75th percentile
Gmyrek et al. (2023)	Gmyrek et al. (2023)	Occupation-level score from the original paper
Gmyrek et al. (2023) - Strict	Gmyrek et al. (2023)	Occupation-level score set to 0 for occupations in the "No Impact" and "Big Unknown" categories
Webb (2019) - 50th pct	Webb (2019)	Exposed occupation: AI score > 50th percentile
Webb (2019) - 75th pct	Webb (2019)	Exposed occupation: AI score > 75th percentile

It is important to stress that these estimates generally remain agnostic about whether AI will serve as a substitute or a complement to human labor. The exception is Gmyrek et al. (2023), using some judgment by the researchers themselves. They postulate that AI would substitute (complement) occupations when the mean task-level exposure is high (low) and the standard deviation across tasks is low (high). Further discussion on the issue of complementarity is provided in Pizzinelli et al. (2023) and Cazzaniga et al. (2024). The distinction between substitution and complementarity, however, does not affect the results of the analysis. As Acemoglu (2024) notes, the increase in TFP derived using Hulten's theorem can arise from both automation (i.e., substitution) or complementarity, as long as it represents cost savings in production.

Appendix 3: AI Adoption Rate

In the framework of Acemoglu (2024), the adoption rate is the proportion of tasks that firms will find profitable to automate through AI, out of all AI-exposed tasks. Even if AI is highly capable, meaning that many tasks are ‘exposed’ to AI automation (as explained in the preceding Appendix), AI will only boost productivity to the extent that it is also cheap enough for firms to adopt. For the United States, Acemoglu (2024) calibrates this adoption rate at 23 percent. This number is the proportion of computer vision tasks that Svanberg et al. (2024) find would be profitable for each firm to automate by deploying, maintaining, and running its own computer vision AI model.⁷

How different would this adoption rate be for Europe as a whole and how might it vary across European countries? To examine this, we use data on AI adoption in Europe, defined as the percentage of firms that use at least one of these technologies: text mining, speech recognition, natural language generation, image recognition/processing, machine learning, software-based workflow automation or decision making, and autonomous machines. Eurostat publishes these data for each country and 2-digit NACE sector (excluding agriculture). They show that around 8 percent of all European firms used at least one type of AI in 2023.

Data are only available on the share of *firms in a particular country and sector* that use AI, whereas Acemoglu’s (2024) framework requires us to calibrate the proportion of *tasks* that will be automated by AI. Our calibration for European countries starts from Acemoglu’s (2024) calibration for the US and adjusts it up or down for each European country based on underlying drivers of AI adoption. We therefore effectively assume that the drivers of AI adoption are similar regardless of whether AI adoption is measured as the proportion of enterprises or the proportion of tasks.

We examine the drivers of AI adoption in the long run using the following cross-section regression:

$$y_{i,j} = \alpha_i + \gamma_j + x_{i,j}\beta + e_{i,j}, \quad (1)$$

where $y_{i,j}$ is the AI adoption rate in country i and sector j , and $x_{i,j}$ are potential drivers of AI adoption including labor and capital costs, industry concentration, and digitalization.⁸ By focusing on the drivers across countries and sectors in equation (1), we effectively focus on fundamental, structural drivers of AI adoption, as opposed to drivers that might temporarily raise or lower the rate of AI adoption relative to an underlying trend.⁹

Appendix Table A2 shows the results of estimating equation (1). Column (1) suggests a strong and robust relationship between AI adoption and wages, where countries and sectors with higher wages tend to have higher rates of AI adoption. Every one euro per hour increase in wages is associated with a 0.75 percentage point increase in the share of firms using AI in each country and sector, controlling for unobserved factors at

⁷ The distinction between developing AI models from scratch or simply fine-tuning existing AI models is not entirely clear from Svanberg et al. (2024). Either way, the authors’ estimate of the cost of the median firm’s AI system is substantial, at between 2 and 3 million US dollars (Figure 5). The authors’ baseline estimate assumes a discount rate of 5 percent a year, an annual reduction in computing costs of 22 percent a year, and a useful life of 5 years for an AI system.

⁸ α_i , γ_j , and β are deterministic parameters to be estimated using ordinary least squares, and $e_{i,j}$ is an error term that we assume to be uncorrelated with the drivers $x_{i,j}$.


⁹ The variables used in the regressions are defined in the footnote to Appendix Table A2 and descriptive statistics for these variables are included in Appendix Table A3.

the country and sector level. This result is consistent with the hypothesis that firms with higher wages have more incentive to use AI to reduce labor costs and/or to increase worker productivity. This result is also robust to excluding country- or sector-specific factors (Column (2)), changing the year in which wages are measured (Column (3)), using labor costs to the firm instead of wages received by the worker (Column (4)), or controlling for other factors (Columns (11)-(12)).

Appendix Table A2: Determinants of AI adoption over the long run.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
wage (2016)	0.75*** (0.22)	0.52*** (0.10)			0.62*** (0.09)	0.40*** (0.07)	0.78*** (0.30)	0.52** (0.22)	0.39*** (0.09)	0.74*** (0.18)	0.22*** (0.05)	0.25*** (0.07)
wage (2021)			0.57** (0.24)									
labor cost				0.58*** (0.17)								
cost of capital					0.35 (0.50)	-0.96** (0.45)					-1.34*** (0.44)	
concentration							-0.03 (0.02)				-0.02 (0.04)	-0.02 (0.05)
computer use								-0.05 (0.06)			0.05 (0.04)	-0.03 (0.05)
energy cost									-6.12 (17.34)		-25.61 (19.70)	
human capital										0.09 (0.11)	0.003 (0.05)	0.01 (0.07)
country FE:	yes	no	yes	yes	no	no	yes	yes	no	yes	no	yes
sector FE:	yes	no	yes	yes	no	yes	yes	yes	yes	yes	yes	yes
# countries	30	30	27	30	25	25	24	24	30	28	16	17
# sectors	11	11	11	11	11	11	11	7	11	10	7	7
# observations	265	265	252	265	226	226	202	153	265	232	102	109

Notes: *p<0.1; **p<0.05; ***p<0.01. Standard errors clustered by country and sector are shown in parentheses. The dependent variable is the proportion of firms within each country--sector using at least one type of AI in 2023. The wage variable measures wages, salaries, bonuses, allowances, employer contributions to saving schemes, and remuneration in kind, as received by workers, in euros per hour. Labor costs measure costs to firms, which include wage and non-wage costs minus subsidies. Cost of capital is the harmonized bank lending rate on outstanding loans with more than 5 years' maturity; it varies across countries only. Concentration is the share of firms that have more than 250 employees. Computer use is the proportion of firms using computers. Energy cost is the price of electricity in euros per kilowatt-hour for firms that use between 2 and 20 thousand megawatt-hours per year, including all taxes and fees; it varies across countries only. Human capital is the proportion of employees that have a tertiary education. FE = fixed effect.



The strong positive relationship between AI adoption and wages has two implications: First, since Europe has lower wages than the US, the rate of AI adoption in the average European country is expected to be lower than in the US. We estimate that yearly wages are around 50 percent higher in the US,¹⁰ which translates into a 5-percentage point lower rate of AI adoption in Europe in the long run.¹¹ Therefore, Acemoglu's (2024) calibration of a 23 percent adoption rate for the US implies an adoption rate of 18 percent for Europe as a whole, given lower wages in Europe.

Second, AI adoption should be higher in European countries with higher wages. Wages in Switzerland, Norway, and Denmark, for example, are around double the EU average, while those in Bulgaria, Romania, and Lithuania are around a quarter of the EU average. These big wage differences imply substantial differences in the ultimate rates of AI adoption across European countries. Appendix Figure A1 shows the estimated adoption

¹⁰ This calculation uses OECD data. For example, in 2022, the GDP weighted average wage was €41,408 in Europe, compared to €73,122 in the US, after converting all national wages to euros at average exchange rates. This implies that European wages were 77 percent of US wages in 2022. Calculating these ratios for all years between 2013 and 2022, and averaging across them, gives the 50 percent that we use in this paper.

¹¹ Given an hourly wage of around €19.2 in the EU27 in 2016, the equivalent US wage should be 50 percent higher at around €28.8 per hour, implying a wage differential of €9.6 per hour. Multiplying this wage differential with the slope coefficient of 0.52 from Table 1 Column (2), we find that AI adoption should be around 5 percentage points lower in Europe than in the US. We use the coefficient 0.52 from Column (2), rather than 0.75 from Column (1), because we would like to calibrate AI adoption rates for all European countries and sectors, even those without data on AI adoption or wages.

rates across European countries in 2023, using the model in Table 1 Column (2), which results in a (value added-weighted) average adoption rate across countries of 13.8 percent. Appendix Figure A1 also shows the estimated adoption rates in the medium term, using two different methods (shifting and rescaling) of ensuring that the average future adoption rate is 18 percent, as calculated above. The shifting method, shown in gray crosses in Appendix Figure A1, shifts all the 2023 estimates up by the same amount ($18 - 13.8$), while the rescaling method, shown in orange triangles, rescales the 2023 estimates by multiplying them by the ratio $18/13.8$. Adoption in Bulgaria, for example, might reach only 8 percent under the shifting method, due to low wages, whereas adoption in Switzerland might end up at 29 percent, due to high wages.¹²

We also consider several other possible determinants of AI adoption rates, but in general do not find robust evidence for them; see Columns (5)-(12) of Table A2. In particular:

- *Cost of capital.* Firms facing higher capital costs might be less inclined to adopt AI because they would find it more expensive to obtain financing to invest in AI systems. Indeed, there is some evidence of this behavior in Column (6), with a negative and statistically significant coefficient on the cost of capital, measured by bank lending rates. However, this relationship relies on the use of sector-specific intercept terms and is not robust to their exclusion (Column (5)). Furthermore, since the available data on the cost of capital vary across countries but not across sectors, it is not possible to include country fixed effects in the specification. Note that Column (1) implicitly controls for any cost-of-capital (and other) effects that vary across countries.
- *Concentration.* Firms operating in countries and sectors with less competition might find it easier to adopt AI because they can more easily pass the costs of doing so on to consumers through higher prices, or alternatively, they might have fewer incentives to adopt AI due to their relatively greater market power. The evidence does not support either of these hypotheses (Column (7)). There is no statistically significant relationship between AI adoption and the concentration of firms across countries and sectors.
- *Computer use / digitalization.* Firms that already have more digital infrastructure in place might find it easier to adopt AI (for example, using existing networks and computing resources) and may find it more profitable to do so because they may have more data to which to apply AI tools. However, the evidence does not support such a relationship, as the number of firms using AI is no higher or lower in countries and sectors where more firms use computers (Column (8)). AI adoption shows some positive relationship with the EU's digitalization indices (instead of computer use), but this can be by construction because some of these indices incorporate AI adoption within them.
- *Energy cost.* Developing sophisticated AI models requires large quantities of energy, which can drive up their costs and make them less profitable to adopt. However, there is no statistically significant relationship between energy costs (measured by country-specific electricity costs) and AI adoption (Column (9)). While energy is a critical input to the development of large language models, it is less relevant as an input to the development of other AI models or to the decision of whether to rent the

¹² Specifically, Bulgaria has an AI adoption rate of 4 percent in 2023. Under the shifting method, we add 4.2 percentage points ($=18 - 13.8$), to get a medium-term future adoption rate of 8.2 percent for Bulgaria. Similarly, Switzerland has an adoption rate of 25.1 percent in 2023, and we shift this up by the same 4.2 percentage points as used for Bulgaria and other countries, giving a medium-term future adoption rate of 29.3 percent.


services of an existing large language model.

- *Human capital.* Firms with access to more educated workers might find it easier to adopt AI, not least because their workers can be expected to be more capable of developing and using AI. We find AI adoption to be higher in countries and sectors with more tertiary-educated workers, as shown by Brey and van der Marel (2024). However, this relationship seems to be proxying for the relationship with wages that we find above, because once we control for wages, the coefficient on human capital becomes statistically insignificant (Column (10)). This suggests that firms with more educated workers tend to adopt AI because their workers have higher wages, which gives these firms stronger incentives to reduce their wage bill or increase their highly productive workers' output, and not because their workers are more capable of developing and using AI tools.

Our above calibrations of AI adoption for Europe are consistent with Acemoglu's (2024) choice of a 23 percent adoption rate for the US. In turn, this 23 percent adoption rate is the proportion of firms that Svanberg et al. (2024) estimate would find it profitable to adopt AI if they each had to deploy, maintain, and run their own AI systems (see above footnote in this subsection). However, Svanberg et al. (2024) point to much higher potential adoption rates (four times higher) if firms do not have to build their own AI models for each task but can instead rent an AI service to perform that task from a provider that serves a narrow set of similar firms (e.g., one AI services provider for each automatable task that motor vehicle manufacturers need to perform, another provider for each task that grocery stores need to perform, another set of providers for airlines, and another for banks). If AI can be offered as a service in this way, then all firms within any given (narrow) industry group can effectively share the costs of building, maintaining, and running their specialized AI software for each task. This sharing effectively reduces the costs to each firm substantially through economies of scale, thus boosting industry-wide adoption rates.

In this scenario of AI-as-a-service, Svanberg et al. (2024) estimate AI adoption rates of around 80 percent for the US in the medium term (as opposed to 23 percent used above). Given lower wages in Europe, an 80 percent US adoption rate would be consistent with a 75 percent adoption rate for Europe as a whole, with variation cross European countries according to their wage levels, as shown in the gray crosses in Figure A2. These adoption rates are around four times higher, for the average country, than those in Appendix Figure A1, which would lead to productivity gains that are also around four times higher, on average.

Appendix Figure A2: Alternative scenario of high economy-wide AI adoption rates:
AI-as-a service
 (Percent of tasks)



Source: Authors' calculations.

Notes: The 2023 wage-based estimate (blue dots) is based on the specification in Column (2) of Appendix Table [2], which implies an average estimated adoption rate of 13.8 percent (across countries and sectors, weighted by value added). The scenario for AI adoption in the future (gray crosses) ensures that the average adoption rate for Europe is 75 percent, by using the formula $(2023 \text{ estimate} - 13.8) + 75$. Therefore, these estimates use the shifting method discussed above. Estimates using the alternative, rescaling method are not shown because they produce adoption rates in excess of 100 percent. The economy-wide estimates shown in this figure are based on sectors B-S excluding sector O, using the NACE Rev. 2 classification.

Appendix Table A3: Descriptive statistics for factors that could explain AI adoption rates

variable	no. countries	no. sectors	no. country-sectors	mean	st dev.	min	max	source
AI adoption	30	11	298	9.9	10.2	0.0	84.7	Eurostat Digital Economy and Society
wage (2016)	37	17	580	17.0	13.1	1.5	70.6	Eurostat
wage (2021)	29	17	455	19.6	13.0	1.7	72.2	Eurostat
labor cost	37	17	581	21.6	16.8	1.8	89.5	Eurostat
cost of capital	27	18	482	2.6	1.0	1.2	6.0	European Central Bank
concentration	26	16	361	30.4	21.4	0.0	90.6	Eurostat Structural Business Statistics
computer use	26	8	201	58.5	24.2	17.3	100.0	Eurostat Digital Economy and Society
energy cost	34	18	605	0.1	0.0	0.1	0.2	Eurostat Energy Statistics
human capital	34	18	529	39.3	21.2	4.2	89.7	Eurostat

Notes: variables are defined in the footnote to Appendix Table 2.

Appendix 4: Labor Cost Savings from AI

Appendix Table 4 summarizes microeconomic estimates of the labor cost savings from AI. Typically, they focus on specific tasks or occupations and use randomized trials where the treatment group is given access to AI tools and the control group does not have access to these tools. Acemoglu (2024) uses the average estimates of a subset of these studies to parameterize labor cost savings required for his simulation.

Appendix Table 4: Microeconomic studies

Paper	Occupations and tasks	Methodology	Results
Brynjolfsson et al. (2023)	Customer Care Agent / customer support	Randomized trial using tailored GPT based conversational assistant as treatment	On average issues resolved per hour increased by 14%, including a 34% improvement for novice and low-skilled workers.
Dell'Acqua et al. (2023)	Individual contributor-level consultants / realistic, complex, and knowledge-intensive tasks	Randomized trials using access to GPT-4 AI or GPT-4 AI with a prompt engineering overview as treatment	43% improvement for consultants below the average performance threshold and 17% enhancement for those above it, relative to their individual scores.
Doshi et al. (2023)	Writers / writing creative outputs	Randomized trial using access to Gen AI tool at varying levels as treatment	Writers utilizing the Gen AI tool experienced improvements in novelty and usefulness of 6.7% and 6.4%, respectively,
Kanazawa et al. (2022)	Low-skilled Taxi Drivers / driving	Quasi-experimental design with fixed effects and instrumental variable analysis studying the usage of AI navigation system that helps drivers find customers when a taxi is cruising	Narrows the productivity gap between high- and low-skilled drivers by 14%
Korinek (2023)	Academic Economists / research	Evaluation of LLMs' usability for research projects in ideation and feedback, writing, background research, data analysis, coding, and mathematical derivations	40% of the research tasks experimented were found to be highly useful.
Noy and Zhang (2023)	Range of white-collar occupations / simple writing tasks	Randomized trial using access to ChatGPT 3.5 as treatment	40% faster completion of the task at hand; 18% improvement in quality scores

Peng et al. (2023)	Software Developers / programming	Randomized trials using access to GitHub Copilot as treatment while tasked with writing a web server in JavaScript	Programmers with access to the AI tool, completed tasks 55.8% faster
Schoenegger et al. (2024)	Data Forecasters i.e. Economists / forecasting quantified data for future events	Randomized trial where treatment group uses frontier LLMs to enhance human judgment compared to the control group with less advanced models	Forecasters utilizing advanced LLMs experienced gained a 24% to 28% improvement in predictions accuracy compared to the control group
Kalliamvakou (2022)	Software Engineers / programming	The research comprises two parts: first, a survey of over 2,000 developers using the SPACE ¹ framework to evaluate perceived productivity	88% of developers reported increased perceived productivity when using GitHub Copilot, according to the survey.
Note: The randomized trial result by Peng et al. (2023), which Kalliamvakou participated in, is also documented by Kalliamvakou (2022).			

Appendix 5: National Occupation Regulation

Regulation could hinder AI adoption and limit the effects of productivity gains on tasks in the regulated occupations compared to the similar tasks in non-regulated occupations. These professions typically have strict regulations and require human oversight by trained and licensed workers for complex decision-making or safety. As a result, there may be an extensive regulatory process before AI can be implemented. Consequently, the integration of AI in these fields can be often slower and more cumbersome, creating a barrier to the full realization of AI's potential benefits.

Under *Directive 2005/36/EC (Professional Qualifications Directive)*, a regulated profession is defined as a profession requiring a formal qualification, such as diplomas, state exams, certificates issued by a competent authority, or registration with a professional body for the access and exercise of the profession due to legal or regulatory rules. The European Commission Regulated Professions Database identifies regulated professions for all countries in our sample, including the EU27 countries, Iceland, Norway, Switzerland, and the UK.¹³ The database groups the regulated professions under generic English names of professions which we use to compile comparative data on the number and types of regulated professions across countries. The database covers 563 distinct generic professions that are regulated in at least one country. The number of regulated professions differs; for instance, medical doctors are regulated in all countries, whereas doormen are regulated only in Italy, Ireland, and Denmark based on the information from the database.

Each country is responsible for maintaining its own list of regulated professions in the Regulated Professions Database, and as per disclaimer, the Commission cannot be held responsible for the accuracy of the

¹³ The records are an archived version of the regulated professions of the UK.

information. We take the information of this database at face value, even though there could be gaps in the database in the sense that it does not capture all regulated professions from that list in all countries. For instance, based on the database, airline pilots are not regulated in all countries, which seems counterintuitive. We therefore conjecture that our information compiled from this database is likely to underestimate the true amount of regulation.

We merge the list of generic names of the regulated professions database with preferred or alternate occupational titles at the ISCO-08 4-digit level; out of 563 regulated professions, 366 were matched. The remaining 197 regulated professions in the database cannot be matched with the ISCO-08 database because their generic names are either too general or too specific. Generic names such as 'air flight profession not elsewhere classified' or 'Balloon pilot' are examples of the difficulties of matching to a specific occupation in the ISCO database.

**Appendix Table 4: Matched Regulated Professions to
International Standard Classification of Occupations, 2008 (ISCO-08) 4 Digit Level**
(Percent of unique 4-digit ISCO code)

Country	Regulated Profession	Country	Regulated Profession	Country	Regulated Profession
CZ	32.1	PT	19.4	CY	13.4
SK	26.1	GB	18.5	FI	12.9
HR	24.5	IS	17.6	IE	12.7
AT	24.0	IT	17.1	NO	12.5
LU	22.4	FR	17.1	SE	11.3
CH	22.4	BE	16.6	MT	11.1
PL	21.5	DK	15.7	EE	11.1
DE	21.5	GR	15.5	LT	7.6
HU	21.0	ES	14.8	LV	7.6
SI	20.6	NL	14.8	BG	6.7
LI	19.9	RO	13.4		

Sources: European Commission Regulated Professions Database; ILOSTAT.

Note: ISCO-08 is a four-level hierarchical classification system with 436 unit groups at the lowest level, represented by a 4-digit code. These unit groups encompass over 3,000 preferred occupational titles, indicating that each unit group covers multiple professions. Each occupational title may have up to 89 alternate titles, all of which are used to match regulated professions.

Appendix 6: EU AI Act

The EU AI Act, proposed by the European Commission in April 2021, addresses the regulation on AI technologies. It aims to create a regulatory framework that supports innovation while establishing a consistent legal structure for AI within the EU, ensuring that AI technologies are developed and used safely, ethically, and with respect for fundamental rights. The Act specifies requirements based on level of risk (implemented through defining high-risk systems) and level of computational power (implemented through a risk threshold of

computational power, which we refer to as ‘capacity cap’). While the legislation will take effect in different stages (for instance the section on high-risk systems will enter into force mostly in August 2026), for simplification purposes, we assume that the EU AI act is in effect already or fully anticipated by firms so that they already abide by the regulation.¹⁴

Regarding risk levels, the AI Act prohibits and regulates unacceptable or high-risk systems while imposing lighter to minimal obligations for limited or minimal-risk applications. The high-risk systems refer to systems posing significant risk to safety in critical infrastructure or safety components of products; assessment of access to training, jobs, essential private and public services, and immigration control; and enforcement or administration of law, justice, and democratic processes. Under the regulation, high-risk systems face stricter requirements and obligations when it comes to the application of AI which could substantially limit or delay the use of AI technologies in certain sectors and occupations that use or work with such high-risk systems.

Additionally, the AI Act defines a risk threshold of general-purpose AI model based on the computational power used to train the models (at 10^{25} floating point operations per second or FLOPS). European Commission (2024) highlights that this is 10 times lower than a similar threshold in the United States that was included in the now revoked *Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence* and that a lower threshold can greatly influence the complexity, efficiency, and effectiveness of AI models trained and used. Additionally, European Commission (2024) notes that some state-of-the-art models already exceed the current threshold of the Act.

Appendix 7: Data Privacy Laws

To examine the potential impact of data privacy laws on AI adoption, we select three industries in which data privacy protection is a particular concern, not least because of their data intensity: Information and Communication (NACE Rev. 2 code J), Financial and Insurance Activities (NACE Rev. 2 code K), and Human Health and Social Work Activities (NACE Rev. 2 code Q). We identify these sectors based on recent studies of the impact that the EU General Data Protection Regulation (GDPR) had on these particular industries: Arcuri (2020) examines the impact on the financial sector, Prasad and Perez (2020) and Chen et al (2022) on digital and technology companies, and Yuan and Li (2019) on the health sector. These industries tend to be more data-intensive and could therefore be more affected by data privacy laws. For instance, Yuan and Li (2019) find that hospitals providing digital health services had to make more costly investments in data protection practices once the GDPR entered into force, which ultimately affected their financial performance. Prasad and Perez (2020) suggest that AI use could be limited by the GDPR, posing a challenge particularly for digital services companies. However, Arcuri (2020) notes that while greater investments in data protection infrastructures are necessary, these may ultimately boost productivity. In these sectors, we multiply the occupation-level AI exposure by 0.5, thus implying that the share of tasks exposed to AI in a given occupation is half that of the same occupation performed in an industry not affected by data privacy laws.

¹⁴ EU AI Act Chapter III Section 1 Article 6 Point 1 will enter in to force in August 2027 and Section 4 will enter into force in August 2025.

References

- Acemoglu, D (2024), "The Simple Macroeconomics of AI", NBER Working Paper 32487, <https://doi.org/10.3386/w32487>.
- Acemoglu, D, and P Restrepo (2018), "The Race Between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment." *American Economic Review*, 108(6): 1488-1542.
- Acemoglu, D, and P Restrepo (2019), "Artificial Intelligence, Automation, and Work," *The Economics of Artificial Intelligence: An Agenda*: 197–236. Edited by A. Agrawal, J. Gans and A. Goldfarb. Chicago, USA: University of Chicago Press.
- Acemoglu, D, and P Restrepo (2022), "Tasks, Automation, and the Rise in US Wage Inequality". *Econometrica*, 90(5), 1973-2016.
- Aghion, P and S Bunel (2024), "AI and Growth: Where Do We Stand." Unpublished manuscript. Available at <https://www.frbsf.org/wp-content/uploads/AI-and-Growth-Aghion-Bunel.pdf>.
- Arcuri, M C (2020), "General Data Protection Regulation (GDPR) Implementation: What Was the Impact on the Market Value of European Financial Institutions?", *Eurasian Journal of Business and Economics*, 13(25): 1-20.
- Bergeaud, A (2024), "The Past, Present and Future of European Productivity", ECB Forum on Central Banking, European Central Bank, https://www.ecb.europa.eu/pub/pdf/sintra/ecb.forumcentbankpub2024_Bergeaud_paper.en.pdf.
- Bonfiglioli, A, R Crinò, G Gancia, and I Papadakis (2025), "Artificial Intelligence and Jobs: Evidence from US Commuting Zones", *Economic Policy*, 40(121): 145-194.
- Bradford, A (2024), "The False Choice between Digital Regulation and Innovation", *SSRN Electronic Journal* 118(2), <https://doi.org/10.2139/ssrn.4753107>.
- Brey, Björn, and Erik van der Marel (2024), "The Role of Human-Capital in Artificial Intelligence Adoption", *Economics Letters* 244: 111949.
- Brollo, F, E Dabla-Norris, R de Mooij, D Garcia-Macia, T Hanappi, L Liu and A Nguyen (2024), "Broadening the Gains from Generative AI: The Role of Fiscal Policies", Staff Discussion Notes 2024/002, International Monetary Fund.
- Brynjolfsson, E, D Li, and L Raymond (2023), "Generative AI at Work", NBER Working Paper 31161, <https://doi.org/10.3386/w31161>.
- Cazzaniga, M, F Jaumotte, L Li, G Melina, A Panton, C Pizzinelli, E Rockall, and M Tavares (2024), "Gen-AI: Artificial Intelligence and the Future of Work", IMF Staff Discussion Notes, International Monetary Fund.
- Cerutti, E., A. Garcia-Pascual, Y. Kido, L. Li, G. Melina, M. Tavares, and P. Wingender (forthcoming), "The Global Impact of AI: Mind the Gap", IMF Working Paper, International Monetary Fund.

Chen, C, C B Frey, and G Presidente (2022), “Privacy Regulation and Firm Performance: Estimating the GDPR Effect Globally”, The Oxford Martin Working Paper Series on Technological and Economic Change 2022-1.

Commission de l’Intelligence Artificielle (2024), “IA : Notre Ambition pour la France”, <https://www.elysee.fr/emmanuel-macron/2024/03/13/25-recommandations-pour-lia-en-france>.

Comunale, M, and A Manera (2024), “The Economic Impacts and the Regulation of AI: A Review of the Academic Literature and Policy Actions”, Working Paper 2024/065, International Monetary Fund.

Dell’Acqua, F, E McFowland, E R Mollick, H Lifshitz-Assaf, K Kellogg, S Rajendran, L Kraymer, F Candelon, and K R Lakhani (2023), “Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality”, Working Paper 24-013, Harvard Business School.

Doshi, A R, and O Hauser (2023), “Generative Artificial Intelligence Enhances Creativity”, SSRN Electronic Journal, <https://dx.doi.org/10.2139/ssrn.4535536>.

Eloundou, T, S Manning, P Mishkin, and D Rock (2024), “GPTs Are GPTs: Labor Market Impact Potential of LLMs”, Science 384, no. 6702: 1306–8, <https://doi.org/10.1126/science.adj0998>.

The White House (2023), “Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence”, <https://www.federalregister.gov/documents/2023/11/01/2023-24283/safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/>.

European Commission (2024), “The Future of European Competitiveness”, https://commission.europa.eu/topics/strengthening-european-competitiveness/eu-competitiveness-looking-ahead_en.

Felten, E, M Raj, and R Seamans (2021), “Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses”, Strategic Management Journal, 42(12), 2195-2217.

Filippucci, F, P Gal, A Leandro, C Jona-Lasinio and G Nicoletti (2024a), “The Impact of Artificial Intelligence on Productivity, Distribution and Growth: Key Mechanisms, Initial Evidence and Policy Challenges”, OECD Artificial Intelligence Papers 15, Paris: OECD Publishing.

Filippucci, F., P. Gal and M. Schief (2024b), “Miracle or Myth? Assessing the Macroeconomic Productivity Gains from Artificial Intelligence”, OECD Artificial Intelligence Papers, No. 29, OECD Publishing, Paris, <https://doi.org/10.1787/b524a072-en>.

Frey, C, E Dabla-Norris, R Hornby, and L D Tyson (2024), “Will AI Make It Easier to Get Productivity Growth We Want?”, World Economic Forum, <https://www.weforum.org/stories/2024/10/ai-artificial-intelligencemake-productivity-growth>.

Gmyrek, P, J Berg and D Bescond (2023), “Generative AI and Jobs: A Global Analysis of Potential Effects on Job Quantity and Quality”, ILO Working Paper 96.

Hatzius, J, J Briggs, D Kodnani, and G Pierdomenico (2023), “The Potentially Large Effects of Artificial Intelligence on Economic”, Global Economics Analyst, Goldman Sachs.

Huang, Y (2024), “The Labor Market Impact of Artificial Intelligence: Evidence from US Regions”, IMF Working Papers, no. 2024/199, International Monetary Fund.

Hulten, C (1978), “Growth Accounting with Intermediate Inputs”, *The Review of Economic Studies* 45(3): 511-518.

IMF (2024), “Accelerating Europe’s Income Convergence through Integration”, *Regional Economic Outlook—Europe*, October.

Kanazawa, K, D Kawaguchi, H Shigeoka and Y Watanabe (2022), “AI, Skill, and Productivity: The Case of Taxi Drivers”, NBER Working Paper 30612, <https://doi.org/10.3386/w30612>.

Kalliamvakou, E (2022), “Research: Quantifying GitHub Copilot’s Impact on Developer Productivity and Happiness”, Github Blog, <https://github.blog/news-insights/research/research-quantifying-github-copilots-impact-on-developer-productivity-and-happiness>.

Korinek, A (2023), “Generative AI for Economic Research: Use Cases and Implications for Economists”, *Journal of Economic Literature*, 61(4): 1281-1317.

Larsen and Küspert (2024), “Regulating General-Purpose AI: Areas of Convergence and Divergence Across the EU and the US”, Brookings, <https://www.brookings.edu/articles/regulating-general-purpose-ai-areas-of-convergence-and-divergence-across-the-eu-and-the-us>.

McKinsey (2023), “The Economic Potential of Generative AI. The Next Productivity Frontier.” Available at the [economic-potential-of-generative-ai-the-next-productivity-frontier.pdf](https://www.mckinsey.com/industries/technology-and-digital/ai-the-next-productivity-frontier)

Nathan, A, J Grimberg, and A Rhodes (2024), “Gen AI: Too Much Spend, Too Little Benefit?”, Goldman Sachs Global Macro Research Report 129.

Noy, S, and W Zhang (2023), “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence”, *Science* 381(6654): 187–92, <https://doi.org/10.1126/science.adh2586>.

OECD (2024), “Deadline Extended to 1 October: Public Consultation on Risk Thresholds for Advanced AI Systems”, OECD.AI, <https://oecd.ai/en/work/seeking-your-views-public-consultation-on-risk-thresholds-for-advanced-ai-systems-deadline-10-september>.

Peng, S, E Kalliamvakou, P Cihon, and M Demirer (2023), “The Impact of AI on Developer Productivity: Evidence from Github Copilot”, arXiv preprint, arXiv:2302.06590.

Pizzinelli, C, A J Panton, M M Tavares, M Cazzaniga, and L Li (2023), “Labor Market Exposure to AI: Cross-Country Differences and Distributional Implications”, Working Paper 23/216, International Monetary Fund.

Prasad, A and D R Perez (2020), “The Effects of GDPR on the Digital Economy: Evidence from the Literature”, *Informatization Policy*, Korea Intelligence and Information Society Promotion Agency, 27(3): 3–18.

Schoenegger, P, P S Park, E Karger, S Trott, and P E Tetlock (2024), “AI-Augmented Predictions: LLM Assistants Improve Human Forecasting Accuracy”, *arXiv preprint*, arXiv:2402.07862.

Si, C, D Yang, and T Hashimoto (2024), “Can LLMs Generate Novel Research Ideas? A Large-Scale Human Study with 100+ NLP Researchers”, arXiv:2409.04109.

Svanberg, M, W Li, M Fleming, B Goehring, and N Thompson (2024), “Beyond AI Exposure: Which Tasks Are Cost-Effective to Automate with Computer Vision?”, *Working Paper*, Massachusetts Institute of Technology.

The Economist (2023), “Europe, a Laggard in AI, Seizes the Lead in its Regulation”, December 10.

Webb, M (2019), “The Impact of Artificial Intelligence on the Labor Market”, *SSRN Electronic Journal*, <https://doi.org/10.2139/ssrn.3482150>.

Wintergerst, R (2024), “Künstliche Intelligenz – Wo Steht Die Deutsche Wirtschaft?”, Bitkom, <https://www.bitkom.org/sites/main/files/2023-09/bitkom-charts-ki-im-unternehmen.pdf>.

Yuan, B, and J Li (2019), “The Policy Effect of the General Data Protection Regulation (GDPR) on the Digital Public Health Sector in the European Union: An Empirical Investigation”, *International Journal of Environmental Research and Public Health*, 16(6): 1-20.



PUBLICATIONS

Artificial Intelligence and Productivity in Europe
Working Paper No. WP/YYYY/067