The early impact of AI on firm-level markups and productivity

Preliminary draft

By Sara Calligaris, Flavio Calvino, Andrea Cancellieri, Francesco Filippucci,
Peter Gal, Luca Marcolin, Oliviero Pallanch, Ella Shachar

The views expressed in this paper are solely those of the authors and should do not be interpreted as those of the Organisation for Economic Co-operation and Development (OECD) or its member countries.

July 2026

Abstract

This paper studies the early effects of AI adoption on firm-level productivity and markups across 24 countries by matching AI-related job postings from Lightcast with firm-level financial data from Orbis over the period 2018–2023. The analysis shows that AI adopters were already larger, more productive and characterised by higher markups before adopting AI. Once these pre-existing differences are controlled for in an event-study framework, AI is found to raise productivity significantly among adopters in highly exposed sectors, such as ICT, and among larger firms, with estimated gains of around 4.3% which tend to build up gradually over several years. Gains in other segments of the economy were found to be not significant. The results also suggest that AI increases markups for high-market-share firms and for firms that are *AI frontrunners* in their sector . Taking into account that not all firms adopt AI in the economy, the estimated magnitudes are consistent with 10-year productivity gains of around 0.9–2.4% in annual MFP in the ICT sector for the countries in the sample, in line with moderately optimistic projections for the aggregate economy. Overall, the results suggest that, while AI can generate sizeable productivity gains, these benefits are so far concentrated among specific firms and sectors, potentially widening productivity gaps and reinforcing competition concerns.

1 Introduction

Artificial intelligence (AI) is widely regarded as a general-purpose technology with the potential to significantly transform economies (Agrawal, Gans and Goldfarb, 2019[1]; Filippucci et al., 2024[2]; Calvino, Haerle and Liu, 2025[3]). This paper empirically studies two questions with strong implications for AI's expected economic benefits. First, to what extent will AI raise productivity and how will it affect productivity gaps between leading and lagging firms? This is critical in the context of slowing productivity growth in many advanced and emerging economies over the past two decades, accompanied by an increasing productivity gap across firms (Andrews, Criscuolo and Gal, 2016[4]; Fernald, Inklaar and Ruzic, 2024[5]). Second, will AI strengthen competition, ensuring that consumers benefit from technological progress and firms face the right incentives to innovate, or will it weaken competition and reinforce the decades-long rise in markups and market concentration (Calligaris et al., 2024[6]; Bajgar et al., 2023[7]; De Loecker, Eeckhout and Unger, 2020[8]; Autor et al., 2020[9]), especially in digital sectors (Calligaris, Criscuolo and Marcolin, 2024[10])? The answers to these two questions are closely intertwined: depending on the competitive environment and on cost pass-through, an efficiency gain may be passed on as lower prices, retained as a wider margin, or used to expand market share.

Ex-ante, the impact of AI on competition across the economy is ambiguous (Athey and Morton, 2025[11]; Filippucci et al., 2024[2]; Calligaris et al., 2026[12]; OECD, 2025[13]). On the one hand, AI may reinforce incumbent advantages, as incumbents may be better positioned in terms of scale and resources to invest in AI. On the other hand, AI may lower barriers to entry, facilitate firm growth, and allow smaller and younger firms to challenge incumbents by expanding their capabilities and reducing the cost of complex tasks (Calligaris et al., 2026[12]). AI's effect on productivity, instead, is likely to be positive on average, although it may vary across firms and sectors and emerge with a lag (Brynjolfsson, Rock and Syverson, 2021[18]). AI adoption itself, the pre-condition for benefiting from AI, may be uneven as it depends on complementary assets such as ICT skills and infrastructure, and on organisational capabilities that are mostly present in large and relatively more productive firms (Calvino and Fontanelli, 2023[14]).

Firm-level evidence on the impact of AI adoption on competition remains relatively scant and mixed. Cross-country survey evidence finds no higher market power among AI adopters, though concentration is positively correlated with AI patenting (Calligaris et al., forthcoming[15]). US firm-level studies find markups rising with AI pricing and concentration rising with firms' initial AI investment (Babina et al., 2024[16]; Adams et al., 2026[17]). Regarding productivity, several studies have estimated firm-level productivity impacts, pointing to heterogeneous benefits. Instrumental-variables studies tend to find that AI adopters are more productive (Czarnitzki, Fernández and Rammer, 2023[18]; Aldasoro et al., 2026[19]), while others document productivity effects that arise with a substantial lag (Brynjolfsson, Rock and Syverson, 2021[20]; McElheran et al., 2025[21]), and country-specific event studies find positive effects on employment and output (Aghion et al., 2025[22]; Pastor-Merino et al., 2026[23]; Ropele and Tagliabracci, 2026[24]). No event study, however, spans multiple countries, and few papers examine markups and productivity jointly. This paper addresses both gaps.

This work contributes to the literature by providing new and broad-based firm-level evidence on the effects of AI adoption. It combines firm-level data with a causal empirical strategy to estimate the impact of AI use across 24 countries, providing the first event-study evidence that spans multiple countries and multiple outcomes. Moreover, this is one of the first papers studying the impact of AI on both firm-level productivity and markups, shedding light on whether early AI adoption improved firm performance, altered competitive conditions, or primarily benefited firms already well positioned to exploit new technologies. Accounting for these two outcomes in the same framework is important: depending on the competitive environment and

on cost pass-through, an efficiency gain may be passed on as lower prices, retained as a wider margin, or used to expand market share, a pattern consistent with the joint rise in firm-level productivity, markups and concentration documented over recent decades (Autor et al., 2020[9]; De Loecker, Eeckhout and Unger, 2020[8]).

The analysis combines job postings data from Lightcast to identify companies that seek AI related skills – a proxy measure for AI adoption – with firm-level financial data from Orbis to derive productivity and markups over 2018–2023 – a period that largely predates the recent boom in Generative AI. The identification relies on a staggered difference-in-differences design that compares firms adopting AI at different times over 2018-2023 and controls for firm-specific pre-existing differences between adopters and non-adopters.

The results show, first, that AI adopters were already larger, more productive, and characterised by higher markups before adoption. Once these pre-existing differences are accounted for in an event-study framework, the estimates do not yield a significant effect of AI adoption, on average for on all firms in the sample. However, interesting heterogeneities emerge: AI adoption is found to raise productivity significantly (by 4.3%) in highly exposed sectors – such as ICT – and among larger firms[1]. This magnitude is consistent with projections of aggregate productivity gains from AI for the ICT sector, which are based on task-level estimates of AI gains (Filippucci, Gal and Schief, 2024[25]). Productivity increases build gradually over time, in line with "J-curve" productivity dynamics following new technologies, given significant upfront adjustments and complementary investment needs.

Our results also show that AI adoption leads to higher markups among firms with relatively larger initial market shares, and in firms that adopt AI in industries where AI adoption is initially rare, which we call "*AI frontrunners*". Significant and positive markup effects are detected also in and professional services.

Overall, the results are informative about the channels through which AI operates regarding productivity and mark-ups. Three margins stand out. Along the first, which depends from adoption of potential competitors in the same industry, markup gains concentrate among *,* i.e. those firms posting AI vacancies in industries with initially low adoption. In this case, higher markups are not accompanied by productivity gains, a pattern consistent with AI generating gains that are not passed-through to consumers and retained as first-mover rents. Second, related to firms scale and market position, productivity gains tend to rise with firm size while markup effects rise with market share. Hence , while larger firms are able to increase more efficiency thanks to AI, in the presence of large market shares firms tend to capture the gains and raise markups. The third margin is sector-specific: AI adoption has significant impact on firms in high-exposure sectors, such as ICT and Professional Services, consistently with the fact that AI may impact a larger share of tasks in these sectors. However, while AI raises productivity in ICT, suggesting the gains could be passed through to consumers there, it increases markups in professional services, suggestings AI-generated gains could be retained by producers through higher prices.

These patterns carry several implications. In downstream markets where firms adopt AI as an input, the cross-sectional average adopter premium appears to reflect selection rather than an advantage conferred by AI, which suggests that generalised competition concerns may be premature, while reinforcing the need of continued monitoring. Possible concerns may instead operate through firm scale and first-mover position, which underscores the need to keep AI-relevant inputs such as compute, data and skills contestable, and to enable broad-based, widespread adoption. Moreover, since we confirm that firms at the productivity frontier are more likely to adopt AI, and that AI adoption may have positive productivity effects in some instances, the rise of AI may *widen* the distance between leading adopters and firms that do not adopt, rather than levelling the playing field. A final note of caution is that our results should be seen as early, initial evidence on these issues: given the rapid pace of AI diffusion and the time its effects may

[1] Appendix Table A11Table A10 shows no evidence that post-adoption productivity effects are stronger in more AI-developer-intensive industries, leaving open the possibility that the ICT gains realise among AI-user firms.

take to materialise, richer and more timely firm-level data will be needed to assess fully the effects of AI adoption on productivity and competition.[2]

The paper is structured as follows. The next sub-section reviews the related literature. Section 2 describes the data construction process, focusing on the identification of AI-related vacancies from job posting, the estimation of firm-level markups, and the matching of job-posting information with Orbis firm-level data. Section 3 presents the definition of AI adoption and the empirical strategy used to identify its causal effects. Section 4 reports the main results on the impact of AI adoption on productivity and markups, together with heterogeneity analyses, robustness checks, and projected aggregate gains. Section 5 concludes discussing the interpretation and implications of the results.

1.1. Related literature

The paper is part of a growing literature discussing the impact of AI on productivity and markups at the microeconomic level. A first body of work suggests that AI adoption is linked with higher firm performance, though with pronounced heterogeneity, and increasingly emphasises the role of complementary capabilities. Recent evidence, sometimes focusing on AI innovators, tends to find a more positive association between AI and productivity (Damioli, Van Roy and Vertesy, 2021[26]; Czarnitzki, Fernández and Rammer, 2023[18]). This association remains evident in cross-country settings using harmonised firm surveys (Calvino and Fontanelli, 2023[14]), although productivity advantages of early AI adopters appear, at least to a certain extent, related to the selection of more competitive firms into AI use (Calvino, Costa and Haerle, 2026[27]).

A stronger link is evident in recent work by Aldasoro et al. (2026[19]), who use matched EIBIS–ORBIS data and instrument EU firms' adoption with the adoption rates of comparable US peers, estimating that AI raises labour productivity by roughly 4%, driven by capital deepening rather than employment displacement, with gains concentrated in medium and large firms and conditional on complementary investment in software, data and skills.This identification rests on an exclusion restriction that is strong and not directly testable, namely that US peers' adoption affects an EU firm's outcomes only through its own adoption.

Like this paper, other works follow the “event study” methodology and exploit staggered adoption settings within a single country, which is more robust to endogenous selection into AI adoption. For instance, Aghion et al. (2025[22]) find, on French firm-level data for 2017 to 2020, that employment and sales rise after adoption, while Ropele and Tagliabracci (2026[24]) find, on Italian survey and administrative data, that adoption raises labour productivity by 4.3%, and profitability as well, shifting employment toward higher-skilled roles. Babina et al. (forthcoming) attribute the subsequent growth in sales and employment of AI

[2] It is also worth noting that this paper focuses on the more downstream layers of the AI production chain, as proxied by AI-related skill demand by firms, rather than the issues concerning the upper layers, such as the markets for hardware or for cloud computing (André et al., 2025[87]; OECD, 2025[88]). The competition implications discussed here therefore concern how AI adoption reshapes market structure and conduct among adopting firms, complementing rather than substituting for analyses of competition in upstream AI markets. This analysis also captures competition effects only through markups, which may not detect potential rises in anti-competitive conducts, entry barriers or switching costs that don’t impact markups. A second caveat is that, according to Shapiro and Yurukoglu (2026[43]), markups are best understood as a valuable "early warning system", as they can help detect changes in industry structure and performance and point to potential issues in competition. At the same time they come with a limitation, as they don’t allow to separate a decline in competition from competition in action, since technologies that lower marginal costs while raising fixed costs raise markups even where consumers ultimately benefit. This limitation is intrinsic to the measure, since in the production approach price and marginal cost are not separately identified (Miller, 2025[42]).

investors to product innovation, concentrated in larger firms, and Pastor-Merino et al. (2026[23]) relate productivity to an adoption measure built by web-scraping firm websites and classifying them with large language models. These event studies offer credible within-country identification, but each rests on a single institutional setting. This paper addresses that gap, estimating the effect of AI adoption across 24 countries in a staggered difference-in-differences design, on an adoption measure built from Lightcast job postings, with a correction for selection into adoption.

A critical aspect is the timing with which productivity returns materialise. The literature has in fact highlighted relevant "J-curve" dynamics in which realised gains lag adoption because they require complementary intangible investment and organisational change (Brynjolfsson, Rock and Syverson, 2021[20]; McElheran et al., 2025[21]; Brynjolfsson, Li and Raymond, 2023[28]). Such J-curve dynamics are consistent with evidence suggesting that firms better positioned to integrate AI in their organisational processes – notably developers – are already realizing stronger productivity returns to AI adoption (Calvino and Fontanelli, 2026[29]).

At the worker level, stronger productivity effects of AI are found, often based on experimental studies that focus on specific tasks and in the context of Generative AI. These studies find generally large and positive effects that are typically largest for less experienced or lower-performing workers (see Calvino, Reijerink and Samek (2025[30]) for a literature review). They range between 14 to 60% in various tasks, such as customer support (Brynjolfsson, Li and Raymond, 2023[28]), software development (Cui et al., 2026[31]) , professional writing (Noy and Zhang, 2023[32]), consulting tasks (Dell'Acqua, Kogut and Perkowski, 2025[33]), and representative-sample labour-market experiments (Haslberger, Gingrich and Bhatia, 2023[34]; Cui et al., 2026[31]).

The aggregate implications of these micro effects remain contested. Acemoglu (2025[35]), embedding such estimates in a task-based framework, concludes via a version of Hulten's theorem that the macroeconomic effect is non-trivial but modest (no more than a 0.66% increase in total factor productivity over ten years) whereas more optimistic assessments anticipate substantially larger gains (Aghion and Bunel, 2024[36]; Filippucci, Gal and Schief, 2024[25]; Bontadini et al., 2026[37]). The differences across these studies arise due to their different readings of the available micro level experimental evidence (discussed above) regarding AI's cost saving potential, as well as on sectoral exposure (Eloundou et al., 2024[38]), the pace of adoption (Bick, Blandin and Deming, 2026[39]) and the extent of reallocation.

While a lively debate emerged about recent competition trends in OECD economies, the literature on the relationship between AI and competition, and markups in particular, is relatively scant. Increasing evidence points to rising markups and concentration (Calligaris et al., 2024[6]; Bajgar et al., 2023[7]), accompanied by falling rates of firm entry and job reallocation (Calvino, Criscuolo and Verlhac, 2020[40]). One may possibly interpret these trends as symptoms of weakening competition, for instance due to excessive regulatory barriers (Covarrubias, Gutiérrez and Philippon, 2019[41]); another possibility is that they reflect competition-in-action (Miller, 2025[42]; Shapiro and Yurukoglu, 2026[43]). Central to the second view is the growing weight of intangible investment. Because these assets require substantial upfront costs and reward scale, they tend to favour larger firms, which have both the resources to invest and the ability to fully benefit of the investment returns (De Ridder, 2024[44]; Bessen et al., 2020[45]; Autor et al., 2020[9]). These features display several analogies with AI technologies and can help guiding early empirical investigation of the phenomena.

Most of the early contributions on AI's impact on market competition focused on the market for foundation models, arguing that AI's technological features are conducive to higher concentration (Marar, 2024[46]; Korinek and Vipra, 2024[47]; Azoulay, Krieger and Nagaraj, 2025[48]; Gans, 2026[49]; Schrepel and Potts, 2025[50]). Gans (2026[49]) distinguishes AI outputs from inputs (data, talent and compute) and identifies a mechanism of increasing returns to scale, whereby data and AI attract users whose data in turn improve model performance and reinforce market power (Farboodi and Veldkamp, 2023[51]). This advantage is compounded by network effects and high switching costs, as the ecosystem of tools built around a given

model raises its value while making migration costly (Agrawal, Gans and Goldfarb, 2019[52]). Korinek and Vipra (2024[47]) complement this description by emphasising large training costs, economies of scale and scope, the growing scarcity of compute, skilled workers and data, and a tendency towards vertical integration, while Azoulay, Krieger and Nagaraj (2025[48]) describe complementary assets as competitive moats.

At the same time, recent OECD work points to signs of dynamism in model development: André et al. (2025) document rapid turnover at the frontier and falling quality-adjusted prices, while cautioning that the concentration of critical inputs and the presence of vertically integrated hyperscalers could erode contestability over time, and Demirer et al. (2025[53]) highlight the potential role of open-source models. The picture that emerges is one in which the market for top-tier foundation models is difficult to enter yet strongly contested among leading players, alongside a more competitive segment of lightweight non-frontier models with lower computational requirements (Korinek and Vipra, 2024[47]).

A possible way to relate AI to competition is to consider the mediating role of productivity. Under standard models of imperfect competition (Jaumandreu, 2025[54]), productivity gains raise market shares, so that uneven adoption widens their dispersion; consistent with this, AI users are found to be more productive and larger, though the differential is substantially attributable to complementary assets rather than to AI itself (Calvino and Fontanelli, 2023[14]), with (Babina et al., 2024[16]) finding that gains accrue disproportionately to large firms in already concentrated industries. The same channel may extend to markups (De Loecker, Eeckhout and Unger, 2020[8]; Autor et al., 2020[9]) and operates alongside returns-to-scale and data mechanisms that tend to favour larger firms (Varian, 2019[55]; Lu, Phillips and Yang, 2024[56]) , even as cloud adoption may disproportionately benefit smaller ones (Caldarola and Fontanelli, 2024[57]). For generative-AI users, diffusion may be pro-competitive, fostering business dynamism where human-AI complementarity prevails (OECD, 2025[58]; Calvino, Reijerink and Samek, 2025[30]), while the countervailing concern is that uneven absorptive capacity widens leader-laggard gaps and entrenches first-mover advantages.

More closely related to this work, which studies the impact of AI on competition from a firm dynamics perspective, the empirical literature has studied the drivers of markups and market power. To the best of the authors' knowledge, apart from Calligaris et al. (forthcoming), no study directly relates AI adoption to firm-level markups. Calligaris et al. (forthcoming) find no evidence that AI users set higher markups using cross-sectional variation from ICT surveys, but do find (non-causal) evidence that firms with AI patenting activities exhibit higher markup growth rates, a trend driven mostly by the ICT sectors; within these sectors, firms associated with AI-related patents display higher markups. This pattern might suggest an effective role of AI in creating market power, but only in those sectors where AI constitutes an output rather than an input. This evidence sits within a longer trajectory in which markups are higher in digital-intensive sectors (Calligaris, Criscuolo and Marcolin, 2024[10]).

Beyond this, the literature has approached markups and market power through several channels. The first is demand-side learning and pricing optimisation: (Adams et al., 2026[17]), using a firm-level measure of AI pricing adoption built from Lightcast job-posting data, document that adopters, disproportionately large, high-productivity firms, subsequently grow faster in sales and markups, which their model attributes to AI's role in learning the demand function and pricing closer to the full-information optimum; their identification rests on long differences and is confined to the United States. A second channel is algorithmic pricing and tacit coordination: (Calvano et al., 2020[59]) show that independent Q-learning algorithms reliably learn supracompetitive prices without communication. (Assad et al., 2024[60]) find that, in Germany's retail gasoline market, algorithmic-pricing adoption raises margins, but only in non-monopoly markets, and in duopolies only when both stations adopt, providing evidence that the competitive consequences are genuinely strategic.

Finally, this paper estimates markups using the production approach of De Loecker and Warzynski (2012[61]), which extends the industry-level framework of Hall (1988[62]) to the firm level by recovering the

markup as the wedge between the output elasticity of a variable input and that input's revenue share, without imposing a demand system or assumptions on how firms compete. Bond et al. (2021[63]) caution that, when output elasticities are estimated from revenue rather than physical quantities, the resulting estimator may be uninformative about the level of markups. De Ridder, Grassi and Morzenti (2026[64]) substantially qualify this concern, showing that revenue-based estimates nonetheless recover the dispersion of markups across firms and their evolution over time, and preserve their correlations with other firm characteristics.

2 Data and measurement

2.1. Lightcast

Lightcast (formerly Burning Glass Technologies) provides one of the largest and most widely used databases of online job vacancies available to researchers and policymakers. The data are collected through web scraping of online job advertisements posted on thousands of job boards, company websites, and recruitment platforms. Following collection, vacancies are deduplicated and processed using natural language processing techniques to extract structured information on occupations, industries, employers, locations, education requirements, and detailed skill requirements, among other information. The resulting dataset offers highly granular and near-real-time information on labour demand that is typically unavailable from traditional labour force surveys or employer surveys. The cross-country coverage is also large, including Australia, New Zealand, Canada, the United States, most countries in Europe, and a substantial number of Latin American countries. Against these advantages, the coverage of the dataset is not uniform across sectors, occupations or countries, as some employers rely more on online recruitment than others.

A distinctive feature of the Lightcast data is the richness of the information contained in vacancy texts: detailed descriptions of tasks, technologies, software, tools, and competencies sought by employers. This makes the data particularly valuable for studying technological change, occupational transformation in the task content of occupations and sectors, and the evolution of skill demand, making it possible to identify emerging trends before they become visible in conventional statistical sources.

A growing academic and policy literature, including at the OECD, has leveraged Lightcast data to study the diffusion of AI, changes in skill demand related to AI across countries and occupations, and AI adoption at the firm level (Acemoglu et al., 2022[65]; Babina et al., 2024[16]; Borgonovi et al., 2023[66]). The underlying premise is that firms adopting AI-related technologies seek workers with AI-related skills or advertise positions involving AI tools, algorithms, machine learning, data science, or related competencies. The present study builds on the same approach.

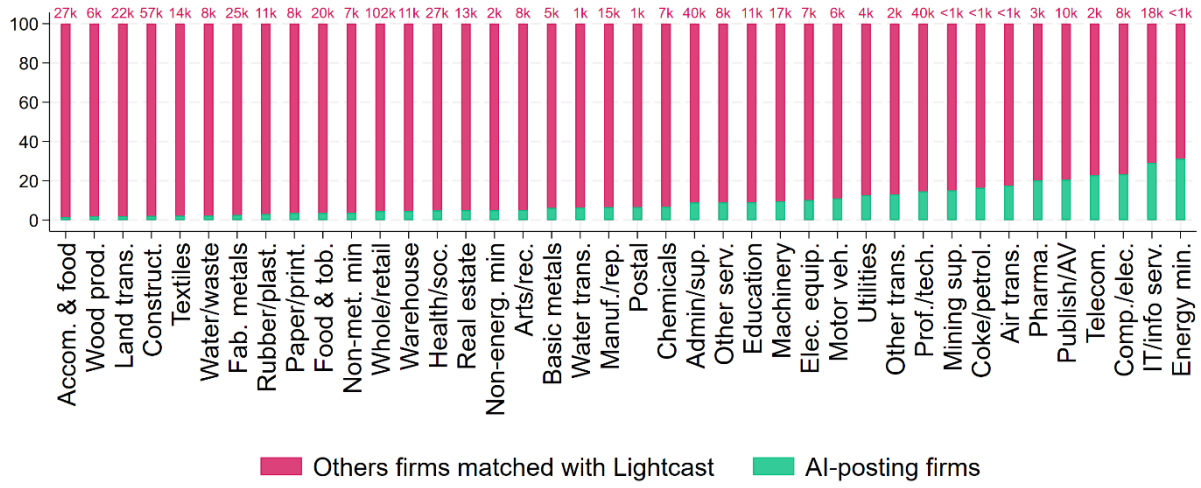
2.1.1. Identifying AI postings based on job ads

The identification of job postings that seek AI related skills follows and updates the approach used by Borgonovi et al. (2023[66]) to identify job vacancies that seek AI related skills. The methodology uses the text of the job ads and classifies them based on whether they include relevant keywords. The set of keywords used by Borgonovi et al. (2023[66]) was expanded to cover developments in AI relevant to our sample peridod, and translate them in the 20 languages that are covered by the Lightcast database in our sample. The resulting set of keywords include more than 900 keywords, classified to be either generic or specific. The requirement of a job ad being classified as seeking AI related skills is that it includes at least either two generic keywords or a specific one.[3]

[3] The method used here relies on specific keywords, which means it might miss AI jobs described in unusual ways or using jargon not in the used keyword list (e.g., saying “cutting-edge algorithms” without saying AI explicitly). Using an expert-validated list should mitigate this risk. The method could also flag a job as AI-related when the context is just mentioning AI in passing. Using the co-occurrence of more than one AI-related keyword to identify AI postings should minimise the frequency of these false positives.

The share of AI intensive job vacancies by sector is similar to what is expected based on AI exposure by sector: knowledge intensive services (ICT services, professional services, publishing and media) rank highly, as well as some manufacturing sectors that are closely related to the production of the technology (computer manufacturing) (Figure 1 and Box 1 below).[4]

Figure 1. Share of firms with AI vacancies by sector



Note: The graph is based on the sample of firms matched between Orbis and Lightcast data, ordered by the share of AI intensive firms. The numbers on top of bars refer to the number of firms present in the matched sample in each sector. See more details on the matching procedure in Subsection 1.1.2 and on the definition of AI intensive firms in Subsection 3.1.
Source: Authors' elaboration on matched Orbis-Lightcast data.

Box 1. Comparing AI exposure measure with an AI job posting measure

This box compares sectoral AI exposure with a job-posting-based indicator of firms' demand for AI related skills, used as an indirect proxy for AI adoption among firms. AI exposure captures the technical potential for AI adoption, while the AI job-posting measure is defined as the share of AI related job postings in total job postings within each industry in the matched Orbis Lightcast dataset. The comparison examines whether industries with higher AI exposure are also those where firms show stronger AI related recruitment demand. This comparison draws on OECD work using online vacancies to measure AI skill demand (Borgonovi et al., 2023[66]), and is also consistent with broader OECD work comparing sectoral indicators of AI related skill demand, AI innovation, AI exposure and actual AI use across industries (Calvino et al., 2024[67]).

The exposure measure is based on the task content of occupations within each industry. Following Eloundou et al. (2024[38]), a task is classified as exposed when an LLM or LLM powered software could reduce the time required to complete the task by at least 50 percent, while maintaining equivalent quality. The industry level exposure value should therefore be interpreted as the share of tasks in the industry that are technically exposed to AI. This makes the measure an indicator of potential AI use at

[4] Some more traditional sectors report high AI incidence (Coke/petrol, mining, transportation), but these are extremely small sectors in our sample (<1k) so the firms observed in these sectors may be a particular selection of highly digital firms actively recruiting online.

the industry level, while actual adoption by firms also depends on technology adoption, worker training, process adjustment and integration into production or service delivery.

For the job posting measure, AI related vacancies are identified within the broader set of online job postings covered by the matched Orbis Lightcast dataset, using structured information on the vacancy, especially the skills required. Firm and industry information in the matched data is then used to aggregate the measure at the industry level. The measure should be interpreted as recruitment based demand for AI related skills, capturing one observable channel through which firms build AI related capabilities. Firms may also adopt AI through existing workers, internal training, external providers or ready-made AI tools. Borgonovi et al. (2023[66]) provide a detailed discussion of the Lightcast data collection and processing procedures, as well as coverage and representativeness issues in online vacancy data.

Figure 2. Sectoral Comparison of AI Exposure and AI Job Posting Rankings, 2022Figure 2 examines whether industries with higher AI exposure also show stronger AI related recruitment demand. It compares industry rankings based on AI exposure with rankings based on AI related job postings. The comparison suggests that the two measures are positively related, but only partially aligned.

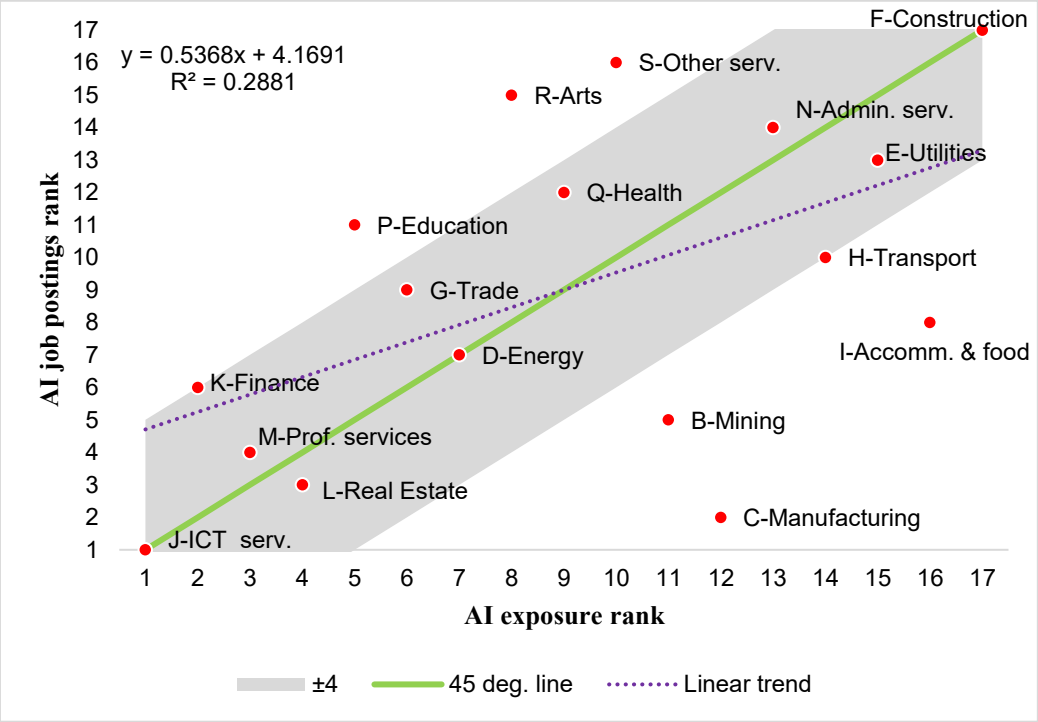
The vertical axis ranks industries by the share of AI related job postings in total job postings within each industry, with rank 1 indicating the highest share. The horizontal axis ranks industries by weighted AI exposure. The fitted trend line has a positive but relatively flat slope: a one rank difference in the AI exposure ranking is associated with only about a half-rank difference in the AI job-posting ranking. This suggests that industries with higher AI exposure tend to show stronger AI related recruitment demand, while the relationship remains partial. This partial alignment is consistent with recent evidence showing that AI exposure and AI adoption are distinct but related dimensions of AI diffusion.[1]

The 45-degree line shows where the two rankings would be identical: industries on this line have the same position in the AI exposure ranking and in the AI job-posting ranking. Because rank 1 denotes the highest value, industries above the line rank lower in AI related job postings than in AI exposure, while industries below the line rank higher in AI related job postings than in AI exposure. Finance, education and health are examples of industries above the line, while mining, manufacturing and accommodation and food services are examples of industries below it.

The positive correlation between AI exposure and the AI job posting measure is also observed in the two previous years, as shown in Appendix A. This persistence over time suggests that the pattern shown in Figure 2 is not driven by a one-year fluctuation in job posting data. Rather, it points to a more stable industry level relationship between potential AI use and firms' demand for AI related skills.

Taken together, the positive correlation supports the relevance of the AI job-posting measure, while the imperfect alignment suggests that industries differ in how far AI exposure is reflected in firms' demand for AI-related skills.

Figure 2. Sectoral Comparison of AI Exposure and AI Job Posting Rankings, 2022



Notes: The vertical axis ranks industries by the share of AI-related job postings in total job postings within each industry. The horizontal axis ranks industries by AI exposure, calculated as a weighted average of country specific exposure values. Weights are based on each country's share of total job postings in the same industry across countries. Rank 1 indicates the highest value. The 45-degree line indicates parity between AI exposure and AI job posting rankings.
Source: Authors'elaboration on Lightcast data.

2.2. Orbis

This paper uses a harmonized firm-level productivity database, based on underlying data from the recently updated OECD-Orbis database (see Gal (2013[68])). The database contains several productivity measures (variants of labour productivity and multi-factor productivity, MFP) and covers 24 OECD countries over the period 1995 to 2024 for the non-farm, non-financial business sector. The underlying data are sourced from annual balance sheet and income statements, collected by Moody's (BVD) – an electronic publishing firm. It is the largest available cross-country company-level database for economic and financial research. The data harmonisation and cleaning steps closely follow Kalemli-Ozcan et al. (2015[69]) and previous OECD work (Bajgar et al., 2020[70]; Andrews, Criscuolo and Gal, 2016[4]).[5]

[5] Three broad steps are: i) ensuring comparability of monetary variables across countries and over time (industry-level PPP conversion and deflation); ii) deriving new variables that will be used in the analysis (capital stock, productivity); and iii) keeping company accounts with valid and relevant information for our present purposes (filtering or cleaning). Finally, Orbis is a subsample of the universe of companies for most countries, retaining the larger and hence probably more productive firms. To mitigate problems arising from this – particularly the under-representation of small firms – the sample is restricted to firms with more than 20 employees on average over their observed lifespan. For more details, see the Annex B.

2.2.1. Estimating firm-level markups

Markups, the ratio of output price, P_{it}, to marginal cost, c_{it}, serve as a firm-level measure of market power. Individual firms' markups are estimated using the production-side approach of De Loecker and Warzynski (2012[61]), which extends the industry-level framework of Hall (1988[62]) to the firm level. The key feature of this approach is that markups can be recovered from the firm's cost minimisation problem conditional on its production technology, without imposing assumptions on the demand structure or the nature of competition. A second advantage of this method is that it is particularly suited to datasets including firms' balance sheets, such as ORBIS, where output prices are not observed and marginal costs have no empirical counterpart.

The method rests on two core assumptions: that at least one input can be adjusted freely within the period, and that firms minimise costs. These are complemented by an assumed functional form for the production function. In this application, the production function is specified as Cobb-Douglas with intermediates serving as the flexible input.[6]

Under these assumptions, the markup can be expressed as the ratio of the output elasticity with respect to the flexible input in a given industry, $\hat{\alpha}_M$, to that input's expenditure share in total revenue, S_{it}^M:

$$\mu_{it} = \frac{P_{it}}{c_{it}} = \frac{\hat{\alpha}_M}{S_{it}^M} \qquad \text{Equation 1}$$

Further details on the derivation of Equation 1 are provided in the appendix. Consistent with standard practice, the baseline production function parameters are assumed to be industry-specific and estimated at the three-digit NACE Rev.2 level, pooling observations across all years and countries in the sample.[7] This choice reflects a trade-off between allowing for heterogeneity and ensuring the precision and stability of the estimated elasticities, given insufficient number of observations to support reliable estimation at a more disaggregated country–industry level. Moreover, for the span of time considered (2018–2023), this empirical strategy appears appropriate, as the relatively short time horizon limits the extent of technological change within industries and supports the assumption that firms operating in the same industry share similar production elasticities. For robustness, also a time varying estimate of elasticities was produced.

It is worth mentioning that the methodology has been scrutinised in a rich debate, with two major conclusions: Bond et al. (2021[63]) point out that, when output elasticities are estimated from revenue rather than physical quantities, the resulting estimator may be uninformative about the *level* of markups. De Ridder, Grassi and Morzenti (2026[64]) reach a more positive assessment, showing that revenue-based estimates nonetheless recover the dispersion of markups across firms, their evolution over time, and preserve their correlations with other firm characteristics. Importantly for this work, the use of fixed effect in regressions (in particular, at least at the industry level) allows to meaningfully use markups in empirical analysis.

[6] The production function has the form $Y_{it} = \Omega_{it} M_{it}^{\alpha_M} L_{it}^{\alpha_L} K_{it}^{\alpha_K}$, with gross output Y_{it}, idiosyncratic productivity Ω_{it}, and three inputs (capital K_{it}, labour L_{it}, and intermediates M_{it})

[7] This level of granularity mitigates a potential issue: competition authorities consider cases within relevant markets, which often go beyond the product-level of disaggregation. Product-level data at the cross-country level and that can be matched with the full array of firm characteristics are rarely available and the three-digit level of disaggregation serves as a valuable proxy.

2.3. Matching

The matching of Lightcast with Orbis is based on company names. Beyond exact matching, a number of steps were taken to increase the matching rate. These include standardising company names (i.e. removing legal status, special characters, etc.) and introducing a fuzzy matching technique to control for typos.[8] Moreover, a supervised learning method is developed based on a small sample to reduce false positive matches that result from fuzzy matching.[9]

The share of firms in the Orbis dataset - with available information on total factor productivity and markups – that were matched with Lightcast based on exact matches is 12% for the median country, which raises to 16% when fuzzy matching is also applied. For several countries where Orbis includes a large number of firms this match rate reaches 60-75%.

[8] Containment matching, Levenshtein Distance matching, No-Suffix matching and Permutation matching were all used to obtain fuzzy matches. When discrepancies emerged across methodologies, the "true" match was chosen using a deep learning model and manual labeling. The model is used to predict the likelihood that a matched pair is a true match, with accuracy being set at 95%.

[9] The matching rate by country is shown in Annex (Figure A1**Error! Reference source not found.**), and the set of countries included in the matching includes AUT, BEL, BGR, CZE, DEU, DNK, ESP, EST, FIN, FRA, GBR, GRC, HUN, IRL, ITA, LVA, NLD, POL, PRT, ROU, SVK, SVN, SWE, USA. Countries are included in the sample if at least 500 firms are matched or the matching rate is at least 5%.

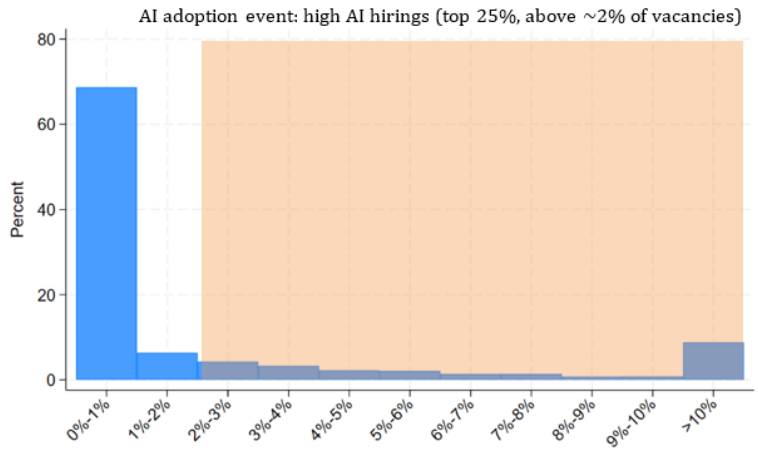
3 Empirical strategy to identify AI's impacts

3.1. Definition of AI adoption

In the data, firms often post only a very small share of AI vacancies. On average, for 75% of firms posting AI-related vacancies in the period 2018-2023, AI-related vacancies account for at most 2% of total vacancies in a year (Figure 3). Such episodes are unlikely to correspond to a significant scale-up in AI adoption by firms. To better identify the effects of AI adoption on firm-level outcomes, this study therefore focuses on firms where AI vacancies represent a more significant share of firms' hiring. Namely, a firm is defined as adopting AI in the first year in which the share of AI-related vacancies over total vacancies exceeds a pre-specified threshold, defined as the 75th percentile (circa 2.2%) of the distribution of AI vacancy shares over the period 2018-2023. From that year onward, the firm is considered as having adopted AI. According to this definition, around 4.63% of matched firms adopt AI in the sample (Table A1).

If firm-level outcomes are available over the period 2018-2023, the use of Lightcast data through 2024 allows identifying adoption until 2024. Table 1 illustrates this setting. Firms that adopt AI in 2018 and firms with only post-treatment observations are excluded, since they are always treated within the estimation sample (1% of matched firms). Firms adopting in 2024 (0.6% of matched firms) are never observed to adopt AI but contribute to the control group as not-yet-treated firms in earlier years.

Figure 3. Distribution of firms with top 25% AI postings



Note: Sample of matched firms that posted at least one AI-related vacancy in the period 2018-2023.
Source: Authors' elaboration on matched Orbis-Lightcast data.

Table 1. Firms adopt AI at different points in time

Illustration of the setting: 1=has adopted AI

Firms	2018	2019	2020	2021	2022	2023
F1 (adopts 2019)	0	1	1	1	1	1
F2 (adopts 2020)	0	0	1	1	1	1
F3 (adopts 2021)	0	0	0	1	1	1
F4 (adopts 2022)	0	0	0	0	1	1
F5 (adopts 2023)	0	0	0	0	0	1
F6 (adopts 2024)	0	0	0	0	0	0
F7 (never adopts)	0	0	0	0	0	0

Note: Each row represents a group of firms (F*) defined by the first year of AI adoption, and each column represents a calendar year. Treatment switches from zero to one in the first adoption year and remain equal to one thereafter.
Source: Authors' elaboration.

3.2. Identification

The empirical analysis is based on the idea that AI adoption may change the trajectory of firm-level outcomes from the year of adoption onward. The data is analysed through the lenses of two models.

The first model is a *between-firm specification* with year fixed effects that compares AI adopters with firms that never adopt AI. It assumes that firm outcomes follow:

$$Y_{it} = \beta_{pre-AI}\ \mathbf{1}(t < G_i) + \beta_{post-AI}\ \mathbf{1}(t \geq G_i) + \tau_t + \theta(\mathrm{X}_i^0 \times \tau_t) + \varepsilon_{it} \qquad \textbf{Equation 2}$$

where Y_{it} is a firm-level outcome, such as log productivity or log markups in calendar year t, G_i is the first year in which firm i adopts AI, and ε_{it} represents random, nonsystemic shocks.[10] The coefficient $\beta_{post-AI}$ captures the average difference in outcome between post-adoption observations of AI adopters and never adopters, conditional on year fixed effects and baseline-characteristic trends.[11] In other words, this coefficient will detect the productivity and markup differential between non-adopters and AI adopters following AI adoption. The coefficient β_{pre-AI} measures instead whether systematic differences between adopters and never adopters were already present *before* AI adoption. The inclusion of year fixed effects, τ_t, absorbs aggregate shocks common to all firms in a given year, while $\mathrm{X}_i^0 \times \tau_t$ controls for differential trends associated with firms' baseline characteristics. Equation 2 is estimated using a linear regression, where controls include trends by firms' baseline characteristics and country-sector-year fixed effects where indicated, which absorb shocks common to firms operating in the same country-sector-year or with similar baseline characteristics.

A second model is a *static event-study specification* (Goodman-Bacon, 2021[71]):

$$Y_{it} = \beta_{post-AI}\mathbf{1}(t \geq G_i) + \gamma_i + \tau_t + \theta(\mathrm{X}_i^0 \times \tau_t) + \varepsilon_{it} \qquad \textbf{Equation 3}$$

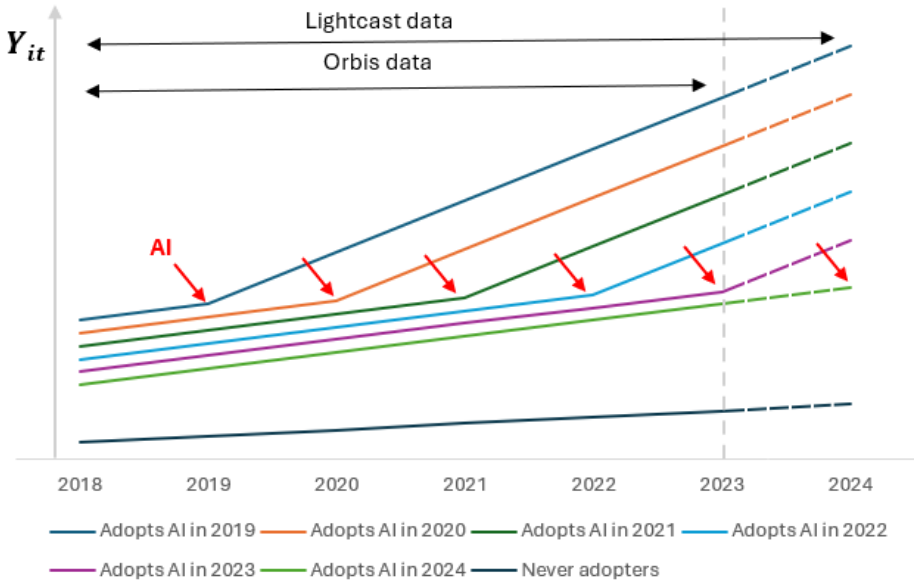
where γ_i are firm fixed effects. Differently from Equation 2, this specification controls for permanent

[10] For never adopters, both event-time indicators (pre and post AI adoption) are set equal to zero, so that they form the reference group.

[11] Baseline-characteristic trends allow firms with different initial size and age to follow different trajectories.

differences across firms through firm fixed effects γ_i. Hence, identification stems from comparing within-firm changes in outcomes following AI adoption to the evolution of the outcome in control firms that have not adopted AI. Control firms may include both firms that, according to Lightcast data, have not yet adopted AI, and firms for which no AI adoption episode is observed (*never adopters)*. However, firms adopting AI are likely to follow structurally different trends in productivity and markups, beyond having different levels of these variables, threatening the identifying parallel-trend assumption. The sample is thus restricted to firms that eventually adopt AI.[12] As a consequence, identification relies on differences in adoption timing rather than on comparison between adopters and never adopters (Figure 4). The coefficient $\beta_{post-AI}$ captures the average post-adoption effect of AI adoption for treated firms, relative to the evolution they would have experienced in the absence of adoption, using firms that have not yet adopted AI as the comparison group, after accounting for permanent firm differences, common year shocks, and differential time patterns by baseline firm characteristics.[13]

Figure 4. Firms adopting AI are compared with firms adopting later



Note: The figure is purely illustrative and does not represent estimated trends. Never adopters are excluded from the estimation sample of Equation 3.
Source: Authors' elaboration.

Equation 3 can be expanded obtaining the following *dynamic event-study specification* (Miller, 2023[72]):

$$Y_{it} = \sum_{e=-5}^{-2} \beta_e \, \mathbf{1}(t - G_i = e) + \sum_{e=0}^{4} \beta_e \, \mathbf{1}(t - G_i = e) + \gamma_i + \tau_t + \theta(\mathrm{X}_i^0 \times \tau_t) + \varepsilon_{it} \qquad \textbf{Equation 4}$$

This specification is very similar to the static event-study but allows to detect the dynamic effects of AI adoption, starting from the year at which firms report high AI postings. Here, $e = t - G_i$ denotes the event time, namely the number of years before or after AI adoption. The year immediately before adoption, $e =$

[12] Excluding firms for which only post-treatment observations are available leaves an estimation sample corresponding to 3.61% of matched firms. See Table A1 in the Annex for details.

[13] As illustrated in 1Part IAnnex C, when using Callaway and Sant'Anna (2021[73]) estimator the effects are relative to contemporaneous changes among firms that have not yet adopted.

-1, is omitted and serves as reference period. The coefficients for $e \geq 0$ capture the evolution of the difference in outcomes after adoption, relative to the reference period. The identifying assumption is that, in the absence of adoption, firms that adopt AI earlier would have followed the same outcome trends as firms that adopt AI later, conditional on fixed effects and controls (conditional parallel-trend assumption). Under this assumption, deviations in the outcomes after the adoption year between adopters and firms that have not yet adopted can be interpreted as the causal effect of AI adoption.

The coefficients for $e < 0$ provide evidence on whether firms that adopted earlier and not-yet-adopters were already on different trajectories before AI adoption. The parallel-trends assumption holds true if these coefficients are not significantly different from zero. However, the identifying assumption would still be violated if firms adopted AI exactly when they were about to experience differential changes in productivity or markups for reasons unrelated to AI. Several features of the empirical design are meant to make sure that the approach partials out any confounding, and that the parallel trend assumption is respected. First, Equation 3 is estimated on the sample of eventual adopters, avoiding direct comparison between adopters and never adopters, that are structurally different. Second, the inclusion of baseline-characteristic trends allows firms with different initial size and age to follow different trajectories. Third, country-sector-year fixed effects absorb shocks in common among firms operating in the same country and sector in a given year.

Equation 3 and Equation 4 are estimated using estimators that are robust to heterogeneous treatment effects (Callaway and Sant'Anna, 2021[73]), allowing treatment effects to differ across adoption cohorts and over time. The estimator uses not-yet-treated firms as control group, so that early adopters are compared with firms that adopt AI later but have not yet adopted in the same calendar year.

3.3. Heterogeneity

In the heterogeneity analysis, the effect of AI adoption is allowed to vary with pre-determined firm-level or country-sector characteristics Z_i^0:

$$Y_{it} = \beta_{post-AI}\mathbf{1}(t \geq G_i) + \beta_{int}\mathbf{1}(t \geq G_i) \times Z_i^0 + \gamma_i + \tau_t + \theta \mathrm{X}_i^0 \times \tau_t + \varepsilon_{it} \qquad \text{Equation 5}$$

Equation 5 is estimated using fixed-effects regressions including firm fixed effects, country-sector-year fixed effects, and baseline-characteristic trends. The interaction coefficient, β_{int}, measures whether the effect of AI adoption is larger or smaller for firms with different initial characteristics, or for firms operating in different country-sector environments.[14] These effects should be interpreted as differential post-adoption changes associated with Z_i^0, relative to pre-treatment outcomes, after accounting for firm fixed effects, common shocks, country-sector-year shocks, and baseline-characteristic trends.

Using this approach, multiple margins of heterogeneity are explored. First, it is analysed whether the effect of AI adoption varies with the baseline intensity of AI-related hiring in the firm's industry, measured as the average share of AI-related vacancies over total vacancies in the industry before adoption. This exercise distinguishes firms adopting AI in sectors where hiring of AI workers was still limited from firms adopting AI that operate in industries where AI-related hiring was more widespread. To ease interpretation, the baseline industry AI-posting measure is mean-centered by subtracting its sample average. With this transformation, $\beta_{post-AI}$ captures post-adoption effect for firms in industries with average baseline AI-posting intensity. The interaction coefficient, β_{int}, captures how this effect changes when the baseline industry AI-posting measure increases by one percentage point. The total effect for a firm i in an industry

[14] In these specifications, Z_i^0 is measured before AI adoption to avoid conditioning on variables that may themselves be affected by treatment. In our setting baseline year is 2018.

j with baseline AI-adoption intensity $Z^0_{j(i)}$ is therefore given by $Effect\left(Z^0_{j(i)}\right) = \beta_{post} + \beta_{int}(Z^0_{j(i)} - \overline{Z^0})$, where $Z^0_{j(i)}$ is the baseline AI-posting intensity of firm i's industry, expressed in percentage points.

Another exercise assesses whether firms adopting AI from a stronger initial market position obtain larger returns from adoption. To construct firm-level market share, real firm-level gross output from Orbis is used as numerator, and real annual gross output at the level of the relevant market from National Accounts as denominator.[15] Following the literature, the logarithm of firm's baseline market share is taken to reduce the influence of the strong skewness in the market-share distribution. This variable is then mean-centred by subtracting its sample average. With this transformation, $\beta_{post-AI}$ captures the effect of AI adoption for a firm with average baseline market share, while β_{int} captures how the treatment effect changes with a one-log-point increase in baseline market share. Equivalently, a 1% higher baseline market share changes the treatment effect by approximately $\beta_{int} \times 0.01$, while a doubling of the baseline market share changes the treatment effect by $\beta_{int} \times \log(2)$. Therefore, the implied treatment effect for a firm with baseline market share Z^0_i is $Effect(Z^0_i) = \beta_{post} + \beta_{int}[\log(Z^0_i) - \overline{\log(Z^0)}]$.

[15] Production data at the 3-digit level of aggregation are obtained combining OECD Structural Business Statistics and Eurostat National Accounts, as detailed in (Calligaris et al., 2026[12]). Relevant markets are defined at the Nace Rev. 2 two-digit level, distinguishing between domestic and internationally tradable markets using the taxonomy defined in Calligaris et al. (2026[12]). For domestic markets, the denominator is computed at the country-sector-year level; for tradable markets, gross output is aggregated across countries within the same sector-year. Due to data limitations United States and Great Britain are excluded from this analysis.

4 Results

	Log MFP		Markup		GO		VA	
	2018	2023	2018	2023	2018	2023	2018	2023
Adopt in 2019	11.98	12.11	1.43	1.54	100.71	142.38	34.60	50.52
Adopt in 2020	11.89	11.88	1.33	1.38	67.40	76.32	24.49	27.72
Adopt in 2021	11.65	11.63	1.31	1.33	53.27	52.95	17.52	17.65
Adopt in 2022	11.68	11.79	1.31	1.37	53.53	55.85	18.81	18.26
Adopt in 2023	11.67	11.78	1.38	1.40	46.22	52.77	15.29	16.95
Never-Adopters	10.85	10.82	1.23	1.22	15.76	14.75	4.27	4.22
	L		LP		K/L		RW	
	2018	2023	2018	2023	2018	2023	2018	2023
Adopt in 2019	300	392	189.03	233.14	54.21	70.56	113.20	119.42
Adopt in 2020	237	311	165.99	234.07	45.19	62.55	95.12	98.22
Adopt in 2021	193	215	135.80	149.53	47.70	54.06	83.69	77.85
Adopt in 2022	206	196	137.19	258.59	44.96	49.22	75.75	82.92
Adopt in 2023	175	193	128.92	150.77	37.81	49.42	79.92	86.47
Never-Adopters	70	70	85.27	90.65	36.43	39.47	51.83	50.05

4.1. Main results: productivity gains in highly AI-exposed country-sectors

Table 2 reports estimates of the relationship between AI adoption and firm-level outcomes. Columns (1) and (2) compare AI adopters with never adopters using the *between-firm specification* in Equation 2. The "Before AI adoption" coefficient captures the average difference in outcomes between future AI adopters, before their adoption year, and firms that never adopt AI. The "After AI adoption" coefficient captures the corresponding difference after adoption.

Both coefficients are estimated conditional on country-sector-year fixed effects. The estimates show that AI adopters are more productive and have higher markups than never adopters, both before and after adoption.[16] This pattern is consistent with previous evidence showing that AI adoption is concentrated among larger, more productive, and more technologically advanced firms (Calvino and Fontanelli, 2026[29]; Babina et al., 2024[16]; Adams et al., 2026[17]).

To provide further evidence on selection into AI adoption

Table A2. Selection into AI adoption I: adopters are larger, more productive, and have higher markups reports average outcomes for AI adopters and never-adopters in 2018 and 2023, while Table A3 estimates the specifications in Columns (1) and (2) for all outcomes, further controlling for trends by baseline firm's characteristics. Compared with non-adopters, AI adopters are larger, more capital intensive, more productive, generate higher value added, pay higher real wages, and display higher markups. This suggests that AI adoption is concentrated among firms with stronger organisational capacity and technological complementarities.

[16] Including baseline-characteristic trends preserves the statistical significance of the coefficients, while reducing their magnitude by about half.

Motivated by these findings, Columns (3) and (4) report the results from running the *static event-study specification* in Equation 3, which controls for any firm-specific difference but based on the sample of eventual AI adopters only. In this setting, the "After AI adoption" coefficient captures the average post-adoption effect of AI adoption for treated firms, relative to the counterfactual evolution of firms that have not yet adopted AI, after accounting for permanent firm differences, common year shocks, and differential time patterns by baseline firm characteristics. The average post-adoption effect is close to zero and statistically insignificant for both productivity and markups. The raw differences between adopters and never adopters in Columns (1) and (2) therefore mainly reflect pre-existing differences across firms rather than the effect of AI adoption itself.

Table 2. Selection into AI adoption, no average effects but heterogeneous productivity effects

	(1)	(2)	(3)	(4)	(5)	(6)
	ln(MFP)	ln(markup)	ln(MFP)	ln(markup)	ln(MFP)	ln(markup)
Before AI adoption	.428***	.037***				
	(.014)	(.006)				
After AI adoption	.491***	.050***	-.002	-.004	.043*	-.023
	(.017)	(.007)	(.015)	(.008)	(.026)	(.017)
Fixed effects						
Firm FEs	N	N	Y	Y	Y	Y
Size, age x year FEs	N	N	Y	Y	Y	Y
Country x sector2 x year FEs	Y	Y	Y	Y	Y	Y
Sample	All	All	AI adopters	AI adopters	High AI exposure	High AI exposure
Observations	652,987	525,138	26,038	23,576	6,541	6,379
Firms	151,921	121,940	5,447	4,944	1,361	1,336
Treated firms	4,382	3,987	3,938	3,566	1,008	980

Notes: AI adoption is defined as the first year in which the share of AI-related vacancies over total vacancies exceeds the 75th percentile of the AI-vacancy-share distribution. Columns (1) and (2) include both AI adopters and never adopters and estimate Equation 2, including country-sector-year fixed effects; the "Before AI adoption" and "After AI adoption" coefficients capture the average difference in outcomes between AI adopters, before and after their adoption year, and firms that never adopt AI. Columns (3)–(6) restrict the sample to AI adopters. Columns (5) and (6) further restrict the sample to observations in the fourth quartile of the AI-exposure distribution. Columns (3), (4), (5), and (6) are estimated using the Callaway and Sant'Anna (2021[73]) estimator and correspond to Equation 3 that includes firm and country-sector-year fixed effects and baseline-characteristic trends. The "After AI adoption" coefficient captures the average post-adoption effect of AI adoption for treated firms, relative to the counterfactual evolution of firms that have not yet adopted AI. Standard errors are clustered at the firm level. Significance levels: * p<0.10, ** p<0.05, *** p<0.01.
Source: Authors' estimation on matched Orbis-Lightcast data.[17]

The non-significant effect in Columns (3) and (4) may occur as the sample concerns a period where AI technologies are still not broadly diffused and limited to pre Gen-AI. Significant effects, therefore, could emerge in sectors that were initially highly exposed to AI, as AI can be applied to a broader set of tasks and deliver higher firm-level gains. Columns (5) and (6) focus on firms operating in countries and sectors that were highly exposed to AI.[18] .

[17] The lower number of observations for the results on markups is explained by missing values on the outcome variable.

[18] AI exposure is measured as the share of production tasks that can be improved by AI in a given country-sector. The survey-based exposure measure was developed by Eloundou et al. (2024[38]) and further elaborated by Filippucci,

Within this group, AI adoption increases MFP by around 4.3% relative to firms that have not yet adopted AI. No statistically significant effect is detected for markups. This suggests that productivity gains from AI adoption are concentrated in environments where the task content of production is more exposed to AI. Moreover, a positive and statistically significant effect is detected on labour productivity (Table A4). Positive effects of AI adoption on labour productivity were found also by Aldasoro et al. (2026[19]), although using an IV strategy that cannot explicitly control for firm-specific unobserved differences, and in event-studies that control for firm-specific effects conducted on Italian (Ropele and Tagliabracci, 2026[24]), Spanish (Pastor-Merino et al., 2026[23]) and German (Czarnitzki, Fernández and Rammer, 2023[18]) firms.[19]

Beyond aggregation biases across sectors of differing initial AI exposure, an alternative explanation for the lack of average productivity effects is the fact that general purpose technologies such as AI require complementary investment in skills, data, software and organisational change that may be poorly measured and may reduce measured productivity initially, before any gain can materialise (Brynjolfsson, Rock and Syverson, 2021[20]). Recent evidence on industrial AI in the United States find J-curve-shaped effects (McElheran et al., 2025[74]), with initial productivity losses followed by later gains. Such a pattern does not appear in the present setting, likely because only early effects of AI adoption can be captured. However, as discussed in Section 4.4, dynamic estimates for the ICT and Industry sector indicate that productivity gains emerge only gradually, becoming, for the ICT sector, statistically significant four years after adoption.

4.2. Heterogeneity results: markup advantages for early-adopters and firms with initially higher market share

The average null effect presented in Table 2 potentially masks heterogeneity across firms. This hypothesis is tested following the approach described in Section 3.3, focusing on two margins that are relevant from a competition perspective, early-adopters vs. late-adopters, and high-market-share firms vs. low-market-share firms.

4.2.1. Sectoral AI frontrunners vs. late-adopters

Columns (1) and (2) in Table 3 examine whether the causal effect of AI adoption varies with the baseline intensity of AI-related postings in the firm's industry. In other words, the analysis distinguishes between the effect of adoption in industries where many firms had already adopted AI at baseline and in industries where adoption was still limited. Namely, The "After AI adoption" coefficient captures the effect of AI adoption for firms operating in industries with average baseline AI-posting intensity, while the interaction coefficient captures how this effect changes when baseline industry AI-posting intensity increases by one percentage point.[20] This analysis is motivated by the idea that the returns to a new technology may depend on how widespread the technology already is among competitors. Early adoption can generate temporary

Gal and Schief (2024[25]), who harmonised the measure across countries. The sample is restricted to firms in the fourth quartile of the AI-exposure distribution ("high exposure").

[19] At the same time, other firm-level evidence finds that AI-investing firms grow faster in sales, but do not experience immediate productivity gains (Babina et al., 2024[16]). The magnitude of the effect in Column (5) found in the ICT sector is close to those in Aldasoro et al. (2026[19]) and Ropele and Tagliabracci (2026[24]) for the aggregate economy.

[20] As described in Section 3.3 the baseline industry AI-posting measure is expressed in percentage points and mean-centered. Estimates account for firm fixed effects, country-sector-year fixed effects and baseline-characteristic trends, including industry AI-posting intensity.

competitive advantage, while these gains may dissipate as adoption becomes more widespread (Gómez and Palomas, 2024[75]).

The results show no statistically significant heterogeneity in productivity effects. For markups, the post-adoption effect for a firm in an industry at mean baseline AI-posting intensity is small and statistically insignificant, around -0.4%. The interaction coefficient is negative and statistically significant: a one-percentage-point higher baseline industry AI-posting intensity reduces the effect of AI adoption on markups by around 0.9 percentage points. This suggests that the markup effect of AI adoption is weaker, and eventually negative, in industries where AI-related hiring was already more widespread before adoption. More concretely, for a firm in the industry with the lowest initial share of AI-related posting the implied effect on markups is positive and statistically significant, at around 2.5%. By contrast, for a firm in the industry with the highest baseline share the effect is negative and statistically significant, at around -6.0%.

The left panel of Figure 5 reports implied effects across percentiles of the baseline AI-posting intensity measure. The effect is positive and statistically significant only at the very bottom of the distribution, while is negative and statistically significant for the top quartile. From a competition perspective, this pattern is consistent with the idea that new digital technologies can create temporary rents and competitive advantage for very AI frontrunners, but that these rents may erode as adoption spreads within the industry. Hence, later adoption may reflect competitive pressure rather than generate additional market-power gains. This interpretation is also consistent with recent cross-sectional evidence from an ECB survey, which suggests that peer pressure plays a role in inducing firms to adopt AI (Chaloupka, Lalinský and Lopez-Garcia, 2026[76]).

4.2.2. High-market-share firms vs. low-market-share firms

The second margin investigated studies whether the effect of AI adoption varies with firms' initial market position, measured by baseline market share[21] (Columns (3) and (4) in Table 3). Indeed, a concern is that AI adoption may strengthen the position of already dominant firms, in line with the evidence that digital technologies and intangible assets can be associated with higher markups, rising concentration and "superstar firm" dynamics (Autor et al., 2020[9]; Calligaris, Criscuolo and Marcolin, 2024[10]; Aghion et al., 2023[77]; Crouzet and Eberly, 2019[78]).
As described in Section 3.3, baseline market share enters the specification in logs and is mean-centred. The "After AI adoption" coefficient therefore captures the effect of AI adoption for firms with average baseline log market share, and the interaction coefficient captures how this effect changes with a one-log-point increase in baseline market share.

Column (3) shows no significant heterogeneity in the effect on productivity. Column (4), instead, points to significant heterogeneity in markup effects. For firms with average baseline log market share, AI adoption is associated with a small reduction in markups of around 0.8%. The positive and significant interaction coefficient indicates that this effect becomes more favourable as baseline market share increases. A one-log-point increase in baseline market share raises the estimated markup effect by around 0.6 percentage points; a doubling of baseline market share increases the effect by around 0.4 percentage points. When evaluated at the lower end of the market-share distribution, the implied effect is negative and statistically significant corresponding to a reduction in markups of around 4.5%. When evaluated at the upper bound of the market-share measure, the implied effect is positive and statistically significant, at around 4.2%.

[21] Market share is proxied by the ratio of firm-level real gross output, from Orbis, to total real gross output in the relevant market, from SBS and Eurostat National Accounts data. Markets are defined at the Nace Rev. 2 two-digit level, distinguishing between domestic and tradable markets: for domestic markets, the denominator is computed at the country-sector-year level; for tradable sectors, it is aggregated across countries within the same sector-year.

The right panel of Figure 5 reports the implied effects across percentiles of the market-share distribution. The effect becomes positive and increasing for firms in the top 10% of the distribution.[22] This is consistent with a "superstar firms" interpretation, whereby new technologies disproportionately benefit firms that are better position to scale up new capabilities. It is also consistent with evidence that leading firms may charge higher prices without necessarily becoming more productive (Gutiérrez and Philippon, 2019[79]; Babina et al., 2024[16]; Adams et al., 2026[17]) and that increases in market power can be concentrated among firms at the top of the distribution (De Loecker, Eeckhout and Unger, 2020[8]).[23]

Table 3. Markup gains concentrate among AI frontrunners and high-market-share firms

	(1)	(2)	(3)	(4)
	ln(MFP)	ln(markup)	ln(MFP)	ln(markup)
AI adopters, after AI adoption	.003	-.004	-.005	-.008**
(effect at the mean)	(.009)	(.004)	(.008)	(.004)
After AI *x* baseline industry	-.008	-.009***		
AI-posting Intensity *(p.p.)*	(.007)	(.003)		
After AI *x* baseline firm's			-.005	.006***
log market share			(.004)	(.002)
Fixed effects				
Firm FEs	Y	Y	Y	Y
Size, age x year FEs	Y	Y	Y	Y
AI-posting share x year FEs	Y	Y	N	N
Country x sector2 x year FEs	Y	Y	Y	Y
Sample	AI adopters	AI adopters	AI adopters	AI adopters
Observations	26,719	24,223	21,177	20,060
Firms	5,594	5,083	4,414	4,195
Treated firms	4,262	3,874	3,384	3,203

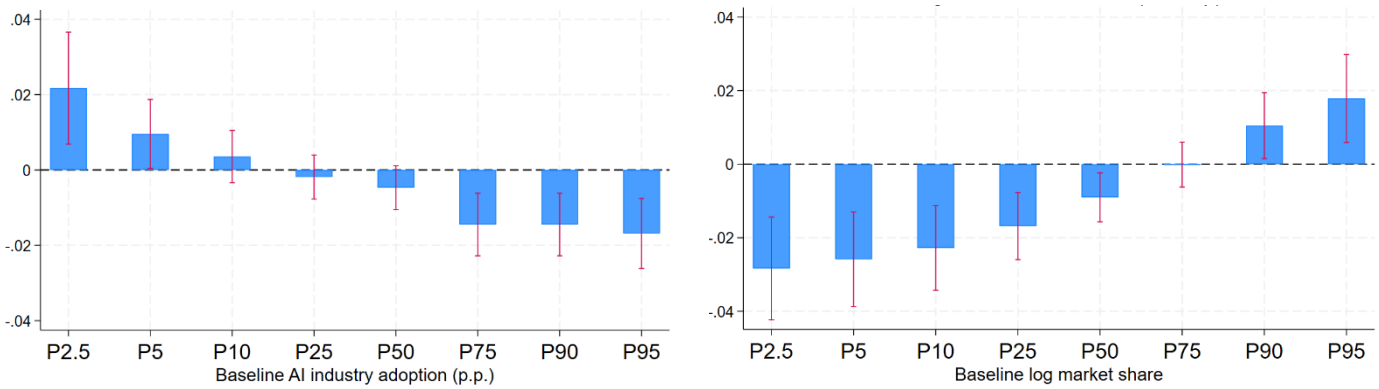
Notes: AI adoption is defined as the first year in which the share of AI-related vacancies over total vacancies exceeds the 75th percentile of the AI-vacancy-share distribution. Columns (1) and (4) include only AI adopters, excluding firms that never adopt AI and estimate Equation 5 that includes firm and country-sector-year fixed effects and baseline-characteristic trends. Columns (1) and (2) further include trends by baseline sector-level AI-posting intensity at the two-digit sector level. In columns (1)-(2), the first coefficient reports the post-adoption effect for firms in industries with average baseline AI-posting intensity, while the second coefficient shows how this effect changes when baseline industry AI-posting intensity increases by one percentage point. In columns (3)-(4), the first coefficient reports the post-adoption effect for firms with average baseline log market share, while the second coefficient shows how this effect changes for a one-log-point higher baseline market share. Standard errors are clustered at the firm level. Significance levels: * p<0.10, ** p<0.05, *** p<0.01.
Source: Authors' estimation on matched Orbis-Lightcast data.

[22] In Appendix, Figure A2**Error! Reference source not found.** shows that firms with high market share are concentrated in the Industry sector while firms with low market share belong primarily to the ICT sector.

[23] For firms starting from weaker competitive positions, the negative effect is more difficult to interpret. One possible explanation is that AI adoption may require complementary investments, such as specialised workers, training, software, and organisational adjustment but these firms may not be able to pass these costs into higher prices (Brynjolfsson, Rock and Syverson, 2021[20]). An alternative explanation is that firms with weaker market positions may use AI adoption to compete more aggressively on prices in the short run to gain market shares. However, these mechanisms cannot be tested with the data at hand.

Figure 5. AI adoption raises markups only for firms with high baseline market share and early AI adopters, while gains fade as AI hiring spreads

Marginal Effect of AI on ln(markup)



Notes: AI adoption is defined as the first year in which the share of AI-related vacancies over total vacancies exceeds the 75th percentile of the AI-vacancy-share distribution. The figure reports implied marginal effects of AI adoption on the logarithm of markup, evaluated at selected percentiles of baseline industry AI-posting intensity (left panel) and at selected percentiles of firms' baseline log market share (right panel). Baseline industry AI-posting intensity is measured at the two-digit sector level, expressed in percentage points, and mean-centred. Estimates are based on Equation 5, using the sample of AI adopters only and excluding firms that never adopt AI. The specification includes firm fixed effects, country-sector-year fixed effects, baseline-characteristic trends including, for the left panel only, trends by baseline sector-level AI-posting intensity. Standard errors are clustered at the firm level. Confidence intervals are reported at the 90% level.
Source: Authors' estimation on matched Orbis-Lightcast data.

Putting the results in Table 2 and Table 3 together suggest that the causal effect of AI adoption is highly heterogeneous. On average, AI adoption does not generate statistically significant productivity or markup gains among adopters. However, productivity gains emerge among firms operating in highly AI-exposed environments, while markup gains are concentrated among firms adopting under more favourable initial conditions: either when AI was still relatively scarce in their industry, or when the firm already had a strong market position.

4.3. Heterogeneity by size

Previous evidence has shown that larger firms are more likely to adopt digital technologies (Calvino, Criscuolo and Ughi, 2024[80]; Forman, 2005[81]) and may be better positioned to benefit from them (Brynjolfsson et al., 2023[82]), potentially as they are more able to exploit them. This motivates Table 4, which investigates whether the causal effect of AI adoption also varies with firm size, measured by firms' baseline employment-size class. The first coefficient reports the post-adoption effect for the reference category, firms with less than 50 employees, while the interaction terms capture additional effects relative to that category.

The estimates show that productivity gains are concentrated among the largest firms with an effect of around 3% for firms with more than 1,000 employees. By contrast, among firms with fewer than 50 employees, AI adoption reduces markups by around 1.9%.[24] These results are in line with evidence showing that the benefits of AI adoption are concentrated among medium and large firms (Aldasoro et al., 2026[19]). Like for past digital technologies, larger firms may be better able to combine AI with complementary assets such as data, digital infrastructure, managerial capabilities and skilled labour. Farboodi et al. (2019[83]) argue that economic activity generates data and large firms generate more data and that induces economies of scale even when the technology of physical production has constant returns, thereby reinforcing scale advantages. Since data is a crucial input to AI implementation, this raises the possibility that AI adoption may reinforce winner-take-most dynamics, rather than uniformly benefiting firms across the size distribution. This interpretation is also consistent with Babina et al. (2024[16]) who find that AI investments are associated with increased industry concentration.

Table 4. AI adoption disproportionately benefits larger firms.

	(1)	(2)
	ln(MFP)	ln(markup)
baseline size: <50 employees	-.019	-.019*
	(.023)	(.011)
additional effect: 50–249 employees	.009	.011
	(.026)	(.013)
additional effect: 250–999 employees	.028	.015
	(.026)	(.013)
additional effect: 1000+ employees	.049*	.022*
	(.029)	(.013)
Fixed effects		
Firm FEs	Y	Y
Size, age x year FEs	Y	Y
Size category x year FEs	Y	Y
Country x sector2 x year FEs	Y	Y
Sample	AI adopters	AI adopters
Observations	26,727	24,231
Firms	5,598	5,087
Treated firms	4,263	3,875

Note: AI adoption is defined as the first year in which the share of AI-related vacancies over total vacancies exceeds the 75th percentile of the AI-vacancy-share distribution. The Table include only AI adopters, excluding firms that never adopt AI and estimate Equation 5, where Z_i^0 is baseline firm-size class. The first coefficient reports post-adoption effect for the reference category (firms with fewer than 50 employees); the interaction terms capture additional effects relative to that category. These effects are measured relative to pre-adoption firm's outcomes, conditional on firm fixed effects, country-sector-year fixed-effects, baseline-characteristic trends, including trends by baseline size category. Standard errors are clustered at the firm level. Significance levels: * p<0.10, ** p<0.05, *** p<0.01.
Source: Authors' estimation on matched Orbis-Lightcast data

[24] For firms with more than 1,000 employees the effect is positive and statistically significant, although small in magnitude (around 0.3%).

4.4. Heterogeneity by sector

Figure A3**Error! Reference source not found.** shows that the share of AI-related vacancies over total vacancies has increased steadily across all sectors since 2018, with important heterogeneity in levels.[25] These patterns suggest that the effect of AI adoption may differ substantially across sectors.

Table 5 explores this heterogeneity. The estimates indicate that AI adoption increases productivity in ICT, with an average post-adoption effect of around 4.3%, relative to the counterfactual evolution of firms that have not yet adopted AI, after accounting for permanent firm differences and common country-sector-year shocks. By contrast, not statistically significant effects on markups are detected in ICT. For Professional Services, the estimates point to a positive and statistically significant effect on markups, of around 3.5%, while the productivity effect is not statistically significant.

Table 5. Productivity effect in ICT, markup effect in Professional Services

	(1)	(2)
	ln(MFP)	ln(markup)
Knowledge-intensive services		
--- ICT (J)	.043*	-.024
	(.025)	(.017)
--- Professional serv. (L-N)	-.022	.035*
	(.029)	(.019)
Industry (B-F)	-.033	-.002
	(.029)	(.011)
Other services	-.049	-.007
	(.038)	(.013)
Fixed effects		
Firm FEs	Y	Y
Country x sector-group x year FEs	Y	Y
Sample	AI adopters	AI adopters
Observations	26,735	24,209
Firms	5,573	5,058
Treated firms	4,045	3,664

Note: AI adoption is defined as the first year in which the share of AI-related vacancies over total vacancies exceeds the 75th percentile of the AI-vacancy-share distribution. Coefficients are estimated using the Callaway and Sant'Anna (2021[73]) estimator and report the average post-adoption coefficient, β_{post}, from Equation 3. Each coefficient is estimated on the subsample of firms belonging to the corresponding sector. These effects are measured relative to the counterfactual evolution of firms that had not yet adopted AI, conditional on firm fixed effects, country-sector-year fixed-effects. Standard errors are clustered at the firm level. Significance levels: * p<0.10, ** p<0.05, *** p<0.01
Source: Authors' estimation on matched Orbis-Lightcast data.

Figure 6 decompose these sectoral effects dynamically by estimating the *dynamic event-study specification* in Equation 4 separately by sector group. This provide a check for the parallel-trend assumption by looking at the coefficients before adoption, that is, before the vertical dashed line. Since these coefficients are close to zero and statistically insignificant, there is no evidence of differentials pre-trends.

[25] ICT remains the sector with the highest average share throughout the period, rising from 1.9% in 2018 to 3.2% in 2024. This is in line with recent evidence showing that AI use is highest in ICT and Professional Services (Calvino and Fontanelli, 2023[14]). Among the other sectors, Professional services and Retail and Wholesale experienced the largest increases, from 0.5% to 0.9%, and from 0.1% to 0.3%, respectively.

Moreover, the Figure offers insights about the dynamic of the effects. In the upper panel, it emerges that the average productivity effect in ICT detected in Table 5 takes time to materialise. Post-adoption coefficients increase gradually and reach around 12.3% four years after adoption. In the Industry sector, instead, there is evidence of short-run productivity losses followed by a later recovery. This pattern is close to what the literature describe as a productivity J curve: productivity may initially decline while the benefits materaialise only later, once AI is fully integrated into production process. This delayed pattern is consistent with the idea that general purpose technologies may require complementary investments and organizational adjustments before productivity gains become visible in measured outcomes, as highlighted by the productivity J-curve literature (Brynjolfsson, Rock and Syverson, 2021[20]).[26] Figure A4 shows that the dynamic effects on labour productivity display a similar pattern to those on productivity, with an effect that reaches 13.8% four years after adoption.

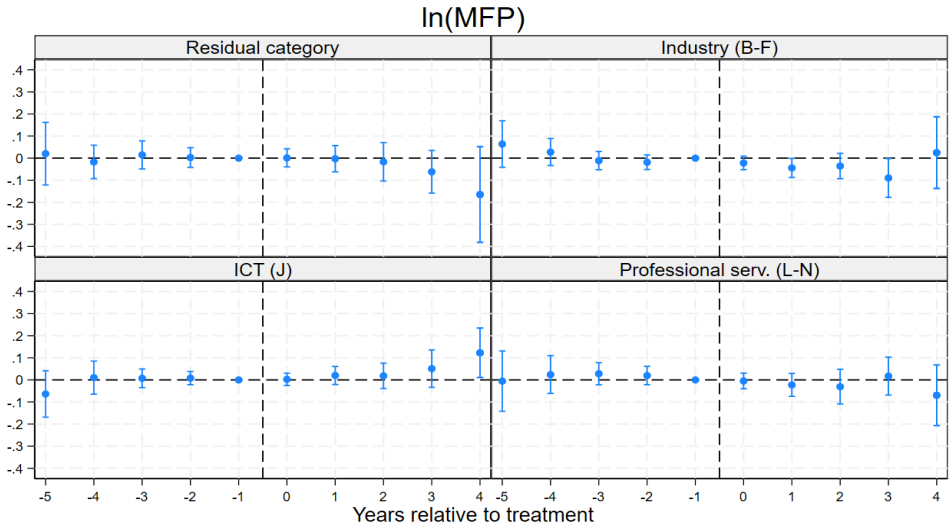
The lower panel focuses on markups. For Professional Services, the markup effect also appears gradually becoming larger over time after adoption and reaching around 10% after four years. In ICT, the initial dip is followed by a later recovery. One possible interpretation is that AI adoption may first require experimentation, organisational adjustment and time to be integrated into services and client relationship before firms are able to translate new AI capabilities into higher margins. However, the analysis does not directly test the underlying mechanism.

Taken together, these results suggest that the effects of AI adoption are concentrated in knowledge-intensive sectors, but along different margins: productivity gains appear stronger in ICT, while markup gains realize in Professional Services. In both cases these effects build gradually over time.

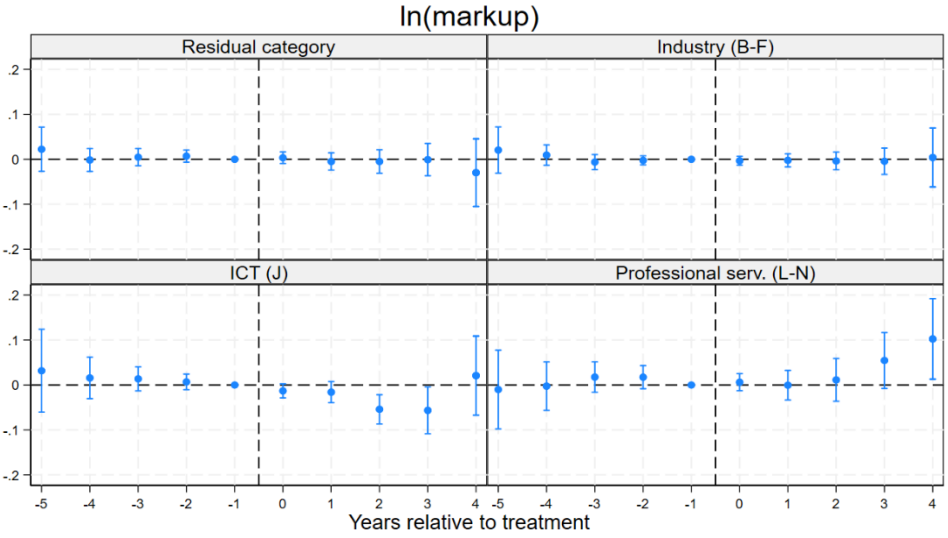
[26] It is possible that post-adoption effects are driven by further increases in AI-related hiring after the adoption threshold is crossed. The evidence points to only mild persistence in AI scale-up over time, suggesting that this channel is unlikely to fully explain the delayed effects.

Figure 6. Dynamic event-study by sector. Productivity effect in ICT, markup effect in Professional Services

Upper panel



Lower panel



Notes: AI adoption is defined as the first year in which the share of AI-related vacancies over total vacancies exceeds the 75th percentile of the AI-vacancy-share distribution. Coefficients are estimated using the Callaway and Sant'Anna (2021[73]) estimator and report the dynamic event-study effects from Equation 4. Each coefficient corresponds to a specific event time relative to the adoption year, with $(t = -1)$ omitted as the reference period. Estimates are computed separately for firms belonging to each sector group and are measured relative to the counterfactual

evolution of firms that had not yet adopted AI, conditional on country-sector-year fixed-effects. Standard errors are clustered at the firm level. Confidence intervals at the 95% level.
Source: Authors' estimation on matched Orbis-Lightcast data.

4.5. Robustness, additional heterogeneity and other outcomes

We conduct several robustness checks to assess whether the results are driven by specific definition of AI vacancies used in the analysis, or by countries and sectors where the AI job-posting measure is not well aligned with survey based AI exposure measure. The results are reported in the Annex (Table A7, Table A8, Table A9). Overall, the signs of the coefficients remain consistent, and statistical significance is generally preserved when restricting the sample to countries and sectors for which the correlation between the AI-posting measure and AI-exposure measure is high.[27] When using the Borgonovi et al. (2023[66]) definition of AI vacancies, the signs of the coefficients also remain consistent although estimates are less precise and statistical significance is weaker.

Further analyses examine whether the estimated effects vary with firms' baseline age, the type of AI activity prevalent in their industry, and their exposure to earlier adopters (Table A11, Table A**12**, and Table A13). Table A11 shows that neither baseline industry AI-user nor AI-developer intensity significantly changes the AI-adoption effect on MFP, whether the two measures are entered separately or jointly. The results do not identify individual firms as AI users or developers. They nevertheless indicate that the productivity response is not systematically stronger in more AI-developer-intensive industries, leaving open the possibility that productivity gains detected in Table 2 and Table 5 extend to AI-user firms too. For markups, both interactions are negative and statistically significant when entered separately, but neither remains significant when included together. Table A**12** reports average post-adoption effects separately across four baseline age groups. The estimates are statistically insignificant in all groups and do not display a clear age gradient. Table A13 examines whether the estimated effects vary with firms' pre-adoption exposure to earlier adopters, measured in the last year in which each firm is observed untreated.[28] The are small and statistically insignificant for both MFP and markups, no evidence of a linear relationship between pre-adoption peer exposure and the effects of AI adoption.[29]

Lastly, we extend the analysis beyond the two main outcomes and examine the effects of AI adoption on profitability.[30] Focusing on this outcome complements the analysis in Table 3 where markup gains are detected for firms adopting AI under more favourable initial conditions. Profitability and markups need not move one-for-one. Markups capture the wedge between prices and marginal costs, while profitability also reflects fixed costs, overhead costs, investment costs, and other accounting components that may differ substantially across sectors. Therefore, observing an increase in profitability alongside an increase in markups would provide stronger evidence for an interpretation based on market power. This is what Column (1) of Table A10 suggests. The interaction coefficient is statistically significant and indicates that firms adopting AI in industries where AI-related hiring was initially lower not only experience higher markups, but also higher profitability. More specifically a one-percentage-point lower baseline share of AI-related postings in the industry increases the estimated effect of AI adoption on profitability by around 2.7

[27] This restriction identifies the part of the data where the job posting proxy provides a clearer signal of AI related firm activity.

[28] Peer exposure is the share of other firms that have already adopted AI in the relevant market-year, with the relevant market defined according to firms' geographic scope. It is measured in the last year in which the firm was observed untreated and peer exposure is measurable. The focal firm is excluded from both the numerator and the denominator.

[29] While this does not rule out spillovers, it suggests that the estimated effects do not display a clear linear gradient with prior AI adoption in the relevant market.

[30] Profitability is measured as the ratio between EBITDA and gross output.

percentage points. By contrast, AI adoption does not appear to have significantly different effects on profitability for firms with high versus low baseline market share.

4.6. Comparison with projected aggregate gains from AI

An important part of the literature on the economic impacts of AI has focused on discussing the potential effects of AI over the next 10 years (Acemoglu, 2025[35]; Filippucci, Gal and Schief, 2024[25]). The debate has often centred on forecasting AI's productivity potential, with radically different ballpark figures. Part of this disagreement reflects different assumptions about how beneficial AI is for the tasks to which it can be applied, and about the share of tasks within firms for which AI could generate productivity gains. The previous section documented significant productivity effects of AI in firms operating in highly exposed sectors, such as ICT. This evidence can be used to take forecasts about the aggregate effects of AI to data: do effects of this magnitude imply meaningful aggregate impacts of AI, or point instead to relatively modest macroeconomic effects?

Following Filippucci, Gal and Schief (2024[25]), and based on Hulten's theorem, the aggregate effects of AI over 10 years for a specific country/industry *cs* can be approximated as the product of task-level performance gains from AI, times how many tasks AI can be applied to in that industry, times the increase in adoption over 10 years in that industry.

$$Aggregate\ effect\ ov\ 10\ years_{cs} \simeq Task\ level\ effect \times exposure_{cs} \times \Delta\ adoption\ (ov\ 10\ y)_{cs}$$

A natural way to link the micro-evidence to the task-based estimates is to assume that the firm-level effect of AI over 10 years can be approximated as the product of the task-level effect and sectoral exposure, namely $Firm\ level\ effect\ (ov\ 10\ y) = Task\ level\ effect \times exposure_{cs}$, as firm-level productivity gains arise from AI being applied to exposed tasks within firms. Under this assumption, the estimates in this paper can be used to recover an implied firm-level effect over a 10-year horizon.

The exercise is limited to the ICT sector, where the analysis in the previous section detects a significant firm-level impact of AI on labour productivity.[31] In this sector, the estimated effect of AI on labour productivity displays a growing dynamic, which is possible to observe only over the first five years in the data available, as shown in Annex Figure A4. The effect over a 10-year horizon may therefore differ depending on how this dynamic evolves beyond the remaining years. Annex Figure A5 describes alternative scenarios for the continuation of the effect over 10 years, and the following exercise will use the firm-level effect of AI adoption over 10 years in the ICT sector resulting from the intermediate scenario.

Not all of the firms are expected to adopt AI over 10 years. To account for this, the firm-level effect of AI adoption over 10 years in the ICT sector is then multiplied by the range of scenarios for the expected increase in AI adoption in the ICT sector over the same horizon, taken from Filippucci et al. (2025[84]), for the 24 countries in the sample. This yields the values reported in the leftmost column of Figure 7, upper panel. Alternatively, it is possible to use the expected adoption for the US only, and obtain the central colum in the Figure. This latter column can be compared to the last column, in red, which is the projected gain for the ICT sector over 10 years according to OECD estimates in 2024 (Filippucci, Gal and Schief, 2024[25]).

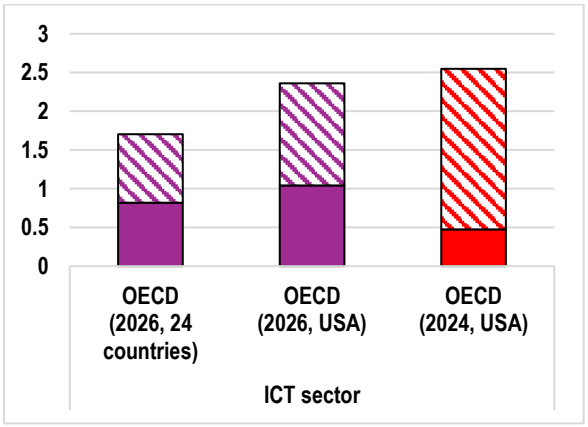
The ICT sector represents a small part of the economy, however, and most of the current macroeconomic debate has focused on the aggregate economy. For benchmarking, one can look at the aggregate estimates implied by Filippucci, Gal and Schief (2024[25]), for the aggregate economy, in the lower panel of Figure 7. Since the micro estimates for the ICT sector in this paper are in line with ICT gains from Filippucci, Gal and Schief (2024[25]), this can be considered as an initial evidence supporting the moderately

[31] This doesn't exclude possible smaller effects in other sectors, where the analysis in this paper may lack power to detect an effect.

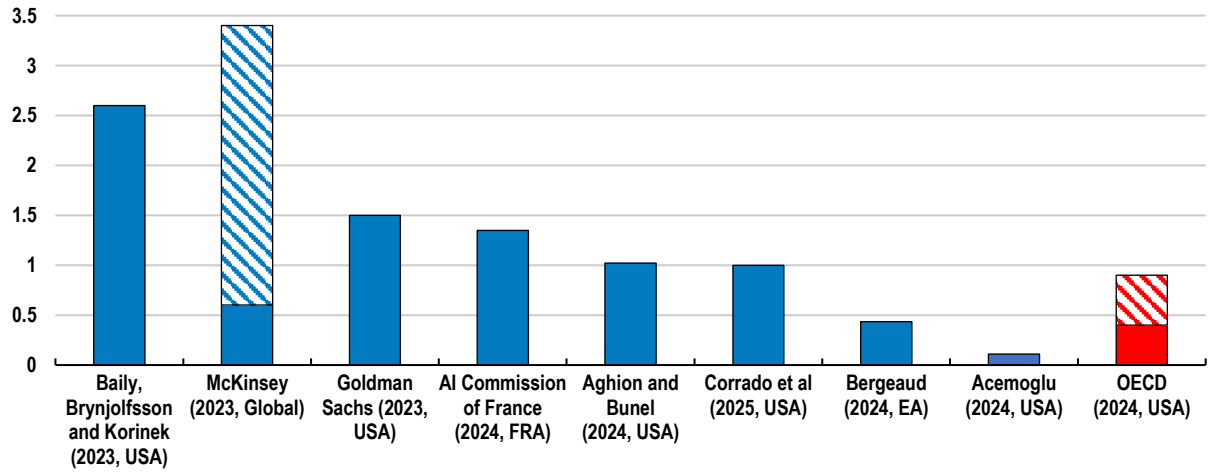
optimistic projections of Filippucci, Gal and Schief (2024[25]) for the aggregate economy. However, it is clearly necessary to keep monitoring gains in other sectors, which may become more detectable as adoption increases.[32]

Figure 7. The estimated firm-level effects in the ICT sector are consistent with past projections of aggregate gains from AI

Upper Panel



Lower Panel



Notes: Purple columns indicate projected aggregate gains derived from the present analysis. Red columns report projections from Filippucci, Gal and Schief (2024[25]). Parentheses report the publication year and the country covered.
Source: Authors' estimation on matched Orbis-Lightcast data, and cited papers.

4.7. Implications for productivity divergence

Past OECD evidence has documented increasing productivity divergence between firms at the global productivity frontier and laggard, or non-frontier, firms, with this divergence being particularly pronounced in digital sectors (Andrews, Criscuolo and Gal, 2016[4]). In this respect, besides detecting productivity

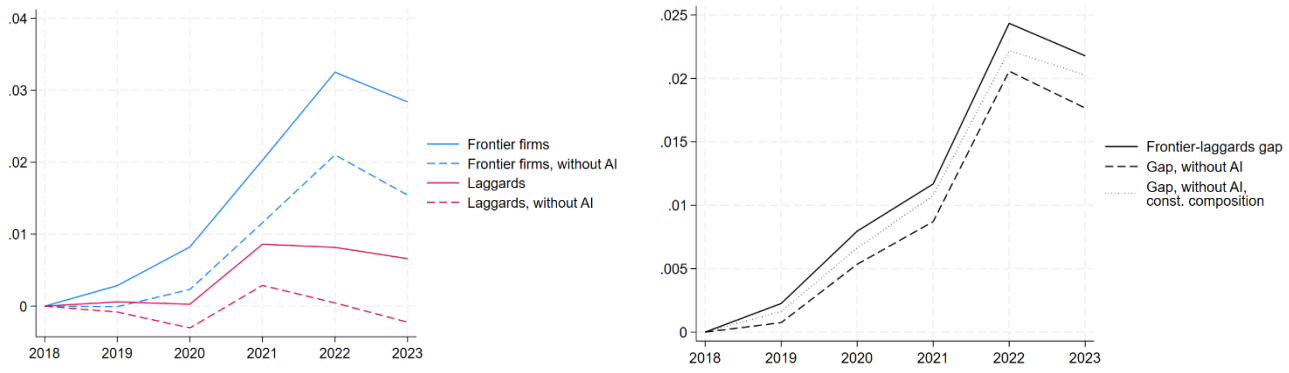
[32] Gains in other sectors may be more affected by J-curve dynamics or be characterised by lower task-level gains. One should also keep in mind that this paper covers mostly pre-Gen AI.

effects in some sectors, the results in this paper show that AI adoption is disproportionately concentrated among high-productivity firms. This raises the concern that AI may contribute to productivity divergence, insofar as its productivity benefits accrue mainly to firms that were already at the productivity frontier before AI adoption, and further benefit in terms of productivity from AI.

Focusing again on the ICT sector, it is possible to estimate how much the combination of higher AI adoption among frontier firms and the productivity benefits associated with AI adoption contributes to frontier-laggard productivity divergence. In the left panel of Figure 8, the solid lines show the path of average MFP for firms in the top decile of the MFP distribution (*frontier firms*), and for the rest of the distribution (*laggards*). These paths imply a widening productivity gap over time, shown by the solid line in the right panel, confirming that the trends observed by Andrews, Criscuolo and Gal (2016[4]) are also present in our sample.

To study the role of AI in this divergence, a counterfactual can be constructed in which frontier firms do not exhibit differentially higher AI adoption and hence do not enjoy AI's productivity benefits more.[33] This counterfactual is shown by the dashed lines in Figure 8. The results suggest that productivity growth among frontier firms would have been roughly halved in the absence of these AI-related channels, pointing to a significant role of AI adoption in the productivity growth of frontier firms in the ICT sector over the period considered. The frontier-laggard gap would also have been around half a percentage point lower, corresponding to roughly one-fifth of the realised gap. This reduction is somewhat smaller, but still present, when comparing frontier and laggard firms that are comparable in terms of sector, size quintile and age quintile, thereby keeping the composition of the two groups constant.

Figure 8. AI has a significant role in the divergence between frontier and laggards firms in terms of productivity in the ICT sector



Notes: Values are normalized to zero in 2018. Frontier firms are firms in the top decile of the productivity distribution; laggard firms are all other firms.
Source: Authors' estimation on matched Orbis-Lightcast data.

[33] This is obtained in two steps. First, differences in predicted AI adoption between frontier firms and other firms are estimated using the following specification: $AI\ adoption_{it} = \gamma_t \times frontier_i \times year_t + \tau_t \times year_t + controls_{it}$, where γ_t and τ_t capture the differential between frontier and laggards in adoption and laggards'adoption, respectively. Second, predicted AI adoption is multiplied by the estimated productivity effect of AI to quantify the contribution of AI. Finally, by subtracting the contribution of AI to the path of MFP for frontiers and laggards it is possible to construct a counterfactual scenario in which the productivity effect associated with AI adoption is set to zero.

5 Discussion and conclusion

This paper contributes to the literature by examining the productivity and competition effects of AI adoption jointly and from a causal perspective, within a single cross-country firm-level framework. Analysing these outcomes in the same framework is important given that the two potentially reinforce one another and are linked in non-trivial ways: an efficiency gain may be passed through as lower prices, retained as a wider margin, or used to expand market share, and only estimating both equations truly allows distinguishing between adoption that raises efficiency and adoption that confers market power.

Three key findings emerge from this analysis. First, there is a considerable degree of selection among AI-adopters, as they differ ex-ante from non-adopters: they are on average larger, more productive, and exhibit higher markups. Second, once these pre-existing differences are accounted for, the average effect of AI adoption on both markups and productivity is not statistically distinguishable from zero. Third, the absence of an average effect on productivity masks substantial heterogeneity across firms and sectors: adoption does in fact affect both outcomes, but its effects vary in sign and magnitude, suggesting that AI adoption effects materialise through different margins. These effects tend to emerge gradually, suggesting a possible J-curve effect, and that effect may become larger over time.

A first margin concerns the conditions under which firms adopt. Markup gains are concentrated among those firms defined as AI frontrunners (in a relative sense, compared to their industry), which are the firms posting AI vacancies in industries with relatively low AI posting intensity at the beginning of the period. No corresponding heterogeneity is found for productivity along this margin, suggesting that what frontrunners may capture is a temporary competitive rent rather than an efficiency premium. This pattern is in fact potentially consistent with a transitory first-mover competitive advantage, which then erodes as adoption spreads in the industry. A risk for competition would emerge if the position of AI frontrunners were to become entrenched. Competition authorities should accordingly monitor whether these rents consolidate in specific settings, in particular through consolidation by leading firms or multiple acquisitions of innovative AI developers that could prevent diffusion from eroding early advantages.

A second margin concerns firm scale and initial market position. Productivity gains rise with firm size, accruing disproportionately to the largest firms while remaining absent among the smallest ones. Markup effects, by contrast, vary more considerably with baseline market shares: they are negative for firms in the lower part of the distribution of market shares, and become positive for firms at the very top of that distribution. The two patterns, while not necessarily coinciding at the firm level, point in the same direction, with the largest and best-positioned firms capturing either productivity gains, or markup gains, or both.

This pattern is in line with broader evidence that digital technologies and intangible assets can reinforce "superstar" dynamics, and raises the possibility that AI adoption strengthens incumbents' positioning rather than levelling the playing field. The appropriate response is to keep AI-relevant inputs contestable, so that the returns to adoption do not accrue disproportionately to incumbents by virtue of scale alone. Access to computational power, data, and skilled labour is central in this respect: where effective AI use depends on a highly skilled workforce and on data assets that scale with firm size, barriers to accessing these inputs may function as barriers to entry. Skills and labour-market policies therefore have direct implications for competition, and supporting broad-based adoption may be a more effective lever than ex-post intervention.

A third margin is sectoral, with relevant heterogeneity emerging when considering sector-specific patterns. Adoption in fact raises productivity in sectors that are highly exposed to AI, most distinctively in ICT, without any accompanying markup increase. Conversely, AI raises markups in professional services, without a measurable productivity gain. The impact in the ICT sector, where efficiency gains are passed through

rather than retained as margin, is consistent with adoption raising productive capacity without conferring market power. Both productivity and markup effects tend to emerge gradually, and the evidence in this paper points to sector-specific J-curve dynamics, as in McElheran et al. (2025[21]). The sector-specific results can also inform the debate about the aggregate impact of AI adoption. Applying Hulten's theorem to the ICT sector estimate and scaling it by expected adoption over a ten years horizon, yields productivity gains that are broadly in line with Filippucci, Gal and Schief (2024[25]). While confined to the ICT sector, the result suggests that the aggregate productivity effects of AI adoption could lie toward the more optimistic end of existing projections.

A further consideration relates to sample composition. Given that a higher share of productivity frontier firms adopt, and adoption leads to a productivity benefit, our results suggest that AI adoption may increase the gap in outcomes between leading adopters and firms that never adopt, rather than acting as a force that levels the playing field. In fact, focusing on the ICT sector, the current analysis has documented that AI appears to have a significant role in the divergence between frontier and laggards firms in terms of productivity.

Overall, the firm-level effects of AI may be better understood as gradual, selective, and sector-specific, and the competition challenges, where they arise, operate through firm scale and first-mover position rather than through adoption as such. At the same time, the productivity gains that emerge in the most exposed sectors and among the largest firms point to a latent dividend that remains unrealised for the wider population of firms. Enabling and facilitating diffusion may help extend these gains beyond the narrow set of sectors and firms in which they currently concentrate, and policies that decrease the cost of adoption and build complementary capabilities are therefore relevant not only for keeping markets contestable but also for ensuring that the productivity benefits of AI are broadly shared.

A note of caution is due, especially with respect to the impact of AI on competition. The analysis in this paper identifies competition effects through firm-level markups. As such, it is worth mentioning a popular perspective (Shapiro and Yurukoglu, 2026[43]), according to which markups are best understood as a valuable "early warning system": they can help detect changes in industry structure and performance and point to potential issues in competition, but they do not allow one to separate a decline in competition from competition in action, since technologies that lower marginal costs while raising fixed costs raise markups even where consumers ultimately benefit. This limitation is intrinsic to the measure, since in the production approach price and marginal cost are not separately identified (Miller, 2025[42]).

Moreover, AI may, affect competition through channels that are not reflected in markups, and, therefore, in the present empirical framework. A first channel operates through entry conditions. Where AI adoption requires large investment in customised solutions, it raises the scale at which adoption becomes profitable; where AI services can instead be bought on demand from cloud providers, that scale is lower and access to data becomes the main constraint. Testing this would require data on firm entry and exit and on the specific type of AI solutions adopted by firms, which the sample used here does not allow. A second channel concerns adopters' reliance on a concentrated set of providers of computing infrastructure and standardised AI solutions. The bundling of these services may generate switching costs and lock-in, weakening competitive pressure even where adopting firms themselves behave competitively. The data do not identify which providers firms rely on, and this dimension is left unexplored. None of these channels is tested in the analysis. They imply, nonetheless, that the absence of a broad-based markup effect should not be read as evidence that AI adoption is competitively neutral.

Finally, this work focuses more closely on the downstream parts of the AI stack, namely the adoption and use of AI by firms. It does not focus on issues arising in the upstream layers of the AI ecosystem, such as chip production, nor does it specifically address competition in the market for foundation models, which deserve separate attention.

More generally, the rapid pace of AI diffusion means that available data do not yet fully capture recent developments, and continued efforts to collect comprehensive information on firm-level adoption would

strengthen the capacity to monitor the evolving diffusion of AI and its impact on the economy. Data that distinguish between AI technology types and actors, capture the intensity and nature of use, and allow adoption events to be tracked over time would be particularly valuable to assess whether the temporary rents documented here dissipate or harden into a structural advantage, and whether productivity gains from AI technologies can spread across firms.

References

Acemoglu, D. (2025), "The simple macroeconomics of AI", *Economic Policy*. [35]

Acemoglu, D. et al. (2022), "Artificial Intelligence and Jobs: Evidence from Online Vacancies", *Journal of Labor Economics*, Vol. 40/S1, pp. S293-S340, https://doi.org/10.1086/718327. [65]

Ackerberg, D., K. Caves and G. Frazer (2015), "Identification Properties of Recent Production Function Estimators", *Econometrica*, Vol. 83/6, pp. 2411-2451, https://doi.org/10.3982/ecta13408. [85]

Adams, J. et al. (2026), "The rise of AI pricing: Trends, driving forces, and implications for firm performance", *Journal of Monetary Economics*, Vol. 157, p. 103875, https://doi.org/10.1016/j.jmoneco.2025.103875. [17]

Aghion, P. et al. (2023), "A Theory of Falling Growth and Rising Rents", *Review of Economic Studies*, Vol. 90/6, pp. 2675-2702, https://doi.org/10.1093/restud/rdad016. [77]

Aghion, P. and S. Bunel (2024), "AI and Growth: Where Do We Stand?", *Policy Note*. [36]

Aghion, P. et al. (2025), "How Different Uses of AI Shape Labor Demand: Evidence from France", *AEA Papers and Proceedings*, Vol. 115, pp. 62-67, https://doi.org/10.1257/pandp.20251047. [22]

Agrawal, A., J. Gans and A. Goldfarb (2019), "Economic Policy for Artificial Intelligence", *Innovation Policy and the Economy*, Vol. 19, pp. 139-159, https://doi.org/10.1086/699935. [1]

Agrawal, A., J. Gans and A. Goldfarb (2019), *The Economics of Artificial Intelligence*, University of Chicago Press, https://doi.org/10.7208/chicago/9780226613475.001.0001. [52]

Aldasoro, I. et al. (2026), "AI adoption, productivity and employment: evidence from European firms", *BIS Working Paper*, Vol. No 1325, https://www.bis.org/publ/work1325.htm. [19]

André, C. et al. (2025), "Developments in Artificial Intelligence markets: New indicators based on model characteristics, prices and providers: New indicators based on model characteristics, prices and providers", *OECD Artificial Intelligence Papers*, No. 37, OECD Publishing, Paris, https://doi.org/10.1787/9302bf46-en. [87]

Andrews, D., C. Criscuolo and P. Gal (2016), "The Best versus the Rest: The Global Productivity Slowdown, Divergence across Firms and the Role of Public Policy", *OECD Productivity Working Papers*, No. 5, OECD Publishing, Paris, https://doi.org/10.1787/63629cc9-en. [4]

Assad, S. et al. (2024), "Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market", *Journal of Political Economy*, Vol. 132/3, pp. 723-771, https://doi.org/10.1086/726906. [60]

Athey, S. and F. Morton (2025), *Artificial Intelligence, Competition, and Welfare*, National Bureau of Economic Research, Cambridge, MA, https://doi.org/10.3386/w34444. [11]

Autor, D. et al. (2020), "The Fall of the Labor Share and the Rise of Superstar Firms*", *The Quarterly Journal of Economics*, Vol. 135/2, pp. 645-709, [9]

https://doi.org/10.1093/qje/qjaa004.

Azoulay, P., J. Krieger and A. Nagaraj (2025), "Old Moats for New Models: Openness, Control, and Competition in Generative Artificial Intelligence", *Entrepreneurship and Innovation Policy and the Economy*, Vol. 4, pp. 7-46, https://doi.org/10.1086/732852. [48]

Babina, T. et al. (2024), "Artificial intelligence, firm growth, and product innovation", *Journal of Financial Economics*, Vol. 151, p. 103745, https://doi.org/10.1016/j.jfineco.2023.103745. [16]

Bajgar, M. et al. (2023), "Industry concentration in Europe and North America", *Industrial and Corporate Change*, https://doi.org/10.1093/icc/dtac059/6987147. [7]

Bajgar, M. et al. (2020), "Coverage and representativeness of Orbis data", *OECD Science, Technology and Industry Working Papers*, No. 2020/06, OECD Publishing, Paris, https://doi.org/10.1787/c7bdaa03-en. [70]

Bessen, J. et al. (2020), "Declining Industrial Disruption", *SSRN Electronic Journal*, https://doi.org/10.2139/ssrn.3682745. [45]

Bick, A., A. Blandin and D. Deming (2026), "The Rapid Adoption of Generative AI", *Management Science*, https://doi.org/10.1287/mnsc.2025.02523. [39]

Bond, S. et al. (2021), "Some unpleasant markup arithmetic: Production function elasticities and their estimation from production data", *Journal of Monetary Economics*, Vol. 121, pp. 1-14, https://doi.org/10.1016/j.jmoneco.2021.05.004. [63]

Bontadini, F. et al. (2026), "AI as an Innovation in the Method of Innovation: Implications for Productivity Growth", *AEA Papers and Proceedings*, Vol. 116, pp. 36-40, https://doi.org/10.1257/pandp.20261036. [37]

Borgonovi, F. et al. (2023), "Emerging trends in AI skill demand across 14 OECD countries", *OECD Artificial Intelligence Papers*, https://doi.org/10.1787/7c691b9a-en. [66]

Brynjolfsson, E. et al. (2023), *Information Technology, Firm Size, and Industrial Concentration*, National Bureau of Economic Research, Cambridge, MA, https://doi.org/10.3386/w31065. [82]

Brynjolfsson, E., D. Li and L. Raymond (2023), "Generative AI at work", *The Quarterly Journal of Economics*, https://academic.oup.com/qje/advance-article-abstract/doi/10.1093/qje/qjae044/7990658. [28]

Brynjolfsson, E., D. Rock and C. Syverson (2021), "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies", *American Economic Journal: Macroeconomics*, Vol. 13/1, pp. 333-372, https://doi.org/10.1257/mac.20180386. [20]

Caldarola, B. and L. Fontanelli (2024), *Cloud Technologies, Firm Growth and Industry Concentration: Evidence from France*, Elsevier BV, https://doi.org/10.2139/ssrn.4965791. [57]

Callaway, B. and P. Sant'Anna (2021), "Difference-in-Differences with multiple time periods", *Journal of Econometrics*, Vol. 225/2, pp. 200-230, https://doi.org/10.1016/j.jeconom.2020.12.001. [73]

Calligaris, S. et al. (forthcoming), *Competition in the age of AI: initial evidence from microdata*. [15]

Calligaris, S. et al. (2024), "Exploring the evolution and the state of competition in the EU", *Protecting competition in a changing world, European Commission.*. [6]

Calligaris, S. et al. (2026), "Defining the geographical level of competition: a taxonomy of industry tradability", *Economica*, Vol. 93/371, pp. 962-991, https://doi.org/10.1111/ecca.70044. [12]

Calligaris, S., C. Criscuolo and L. Marcolin (2024), "Mark-ups in the digital era", *CEP Discussion Papers 1994*, https://cep.lse.ac.uk/_NEW/publications/abstract.asp?index=10858. [10]

Calvano, E. et al. (2020), "Artificial Intelligence, Algorithmic Pricing, and Collusion", *American Economic Review*, Vol. 110/10, pp. 3267-3297, https://doi.org/10.1257/aer.20190623. [59]

Calvino, F., H. Costa and D. Haerle (2026), "Digital technology diffusion in the age of AI: Cross-country evidence from microdata", *OECD Science, Technology and Industry Working Papers*, No. 2026/01, OECD Publishing, Paris, https://doi.org/10.1787/ebc2debe-en. [27]

Calvino, F., C. Criscuolo and A. Ughi (2024), "Digital adoption during COVID-19: Cross-country evidence from microdata", *OECD Science, Technology and Industry Working Papers*, No. 2024/03, OECD Publishing, Paris, https://doi.org/10.1787/f63ca261-en. [80]

Calvino, F., C. Criscuolo and R. Verlhac (2020), "Declining business dynamism: Structural and policy determinants", *OECD Science, Technology and Industry Policy Papers*, No. 94, OECD Publishing, Paris, https://doi.org/10.1787/77b92072-en. [40]

Calvino, F. et al. (2024), "A sectoral taxonomy of AI intensity", *OECD Artificial Intelligence Papers*, No. 30, OECD Publishing, Paris, https://doi.org/10.1787/1f6377b5-en. [67]

Calvino, F. and L. Fontanelli (2026), "AI users are not all alike: The characteristics of French firms buying and developing AI", *Research Policy*, Vol. 55/5, p. 105473, https://doi.org/10.1016/j.respol.2026.105473. [29]

Calvino, F. and L. Fontanelli (2023), "A portrait of AI adopters across countries: Firm characteristics, assets' complementarities and productivity", *OECD Science, Technology and Industry Working Papers*, No. 2023/02, OECD Publishing, Paris, https://doi.org/10.1787/0fb79bb9-en. [14]

Calvino, F., D. Haerle and S. Liu (2025), "Is generative AI a General Purpose Technology?: Implications for productivity and policy", *OECD Artificial Intelligence Papers*, No. 40, OECD Publishing, Paris, https://doi.org/10.1787/704e2d12-en. [3]

Calvino, F., J. Reijerink and L. Samek (2025), "The effects of generative AI on productivity, innovation and entrepreneurship", *OECD Artificial Intelligence Papers*, No. 39, OECD Publishing, Paris, https://doi.org/10.1787/b21df222-en. [30]

Chaloupka, D., T. Lalinský and P. Lopez-Garcia (2026), "What separates firms that use AI intensively from firms that don't?", *The ECB Blog, European Central Bank*, https://www.ecb.europa.eu/press/blog/date/2026/html/ecb.blog20260624~44f70da110.en. [76]

Correia, S. (2017), "Linear Models with High-Dimensional Fixed Effects: An Efficient and Feasible Estimator", *Working Paper*, http://scorreia.com/research/hdfe.pdf. [86]

Covarrubias, M., G. Gutiérrez and T. Philippon (2019), *From Good to Bad Concentration? U.S. Industries over the past 30 years*, National Bureau of Economic Research, Cambridge, MA, https://doi.org/10.3386/w25983. [41]

Crouzet, N. and J. Eberly (2019), "Understanding Weak Capital Investment: the Role of Market Concentration and Intangibles", *NBER Working Paper*, Vol. 25869, https://doi.org/10.3386/w25869. [78]

Cui, K. et al. (2026), "The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers", *Management Science*, https://doi.org/10.1287/mnsc.2025.00535. [31]

Czarnitzki, D., G. Fernández and C. Rammer (2023), "Artificial intelligence and firm-level productivity", *Journal of Economic Behavior & Organization*, Vol. 211, pp. 188-205, https://doi.org/10.1016/j.jebo.2023.05.008. [18]

Damioli, G., V. Van Roy and D. Vertesy (2021), "The impact of artificial intelligence on labor productivity", *Eurasian Business Review*, Vol. 11/1, pp. 1-25, https://doi.org/10.1007/s40821-020-00172-8. [26]

De Loecker, J., J. Eeckhout and G. Unger (2020), "The Rise of Market Power and the Macroeconomic Implications*", *The Quarterly Journal of Economics*, Vol. 135/2, pp. 561-644, https://doi.org/10.1093/qje/qjz041. [8]

De Loecker, J. and F. Warzynski (2012), "Markups and Firm-Level Export Status", *American Economic Review*, Vol. 102/6, pp. 2437-2471, https://doi.org/10.1257/aer.102.6.2437. [61]

De Ridder, M. (2024), "Market Power and Innovation in the Intangible Economy", *American Economic Review*, Vol. 114/1, pp. 199-251, https://doi.org/10.1257/aer.20201079. [44]

De Ridder, M., B. Grassi and G. Morzenti (2026), "The Hitchhiker's Guide to Markup Estimation: Assessing Estimates From Financial Data", *Econometrica*, Vol. 94/1, pp. 137-168, https://doi.org/10.3982/ecta22733. [64]

Dell'Acqua, F., B. Kogut and P. Perkowski (2025), "Super Mario Meets AI: Experimental Effects of Automation and Skills on Team Performance and Coordination", *Review of Economics and Statistics*, Vol. 107/4, pp. 951-966, https://doi.org/10.1162/rest_a_01328. [33]

Demirer, M. et al. (2025), *The Emerging Market for Intelligence: Pricing, Supply, and Demand for LLMs*. [53]

Eloundou, T. et al. (2024), "GPTs are GPTs: Labor market impact potential of LLMs", *Science*, Vol. 384/6702, pp. 1306-1308, https://doi.org/10.1126/science.adj0998. [38]

Farboodi, M. et al. (2019), "Big Data and Firm Dynamics", *AEA Papers and Proceedings*, Vol. 109, pp. 38-42, https://doi.org/10.1257/pandp.20191001. [83]

Farboodi, M. and L. Veldkamp (2023), "Data and Markets", *Annual Review of Economics*, Vol. 15/1, pp. 23-40, https://doi.org/10.1146/annurev-economics-082322-023244. [51]

Fernald, J., R. Inklaar and D. Ruzic (2024), "The Productivity Slowdown in Advanced Economies: Common Shocks or Common Trends?", *Review of Income and Wealth*, Vol. 71/1, https://doi.org/10.1111/roiw.12690. [5]

Filippucci, F. et al. (2024), "The Impact of Artificial Intelligence on Productivity, Distribution and Growth: Key Mechanisms, Initial Evidence and Policy Challenges", *OECD Artificial Intelligence Papers*, OECD Publishing, Paris,, https://doi.org/10.1787/8d900037-en. [2]

Filippucci, F. et al. (2025), "AI meets trade: Global linkages and the cross-country distribution [84]

of the gains from AI", *Forthcoming*.

Filippucci, F., P. Gal and M. Schief (2024), "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. [25]

Forman, C. (2005), "The Corporate Digital Divide: Determinants of Internet Adoption", *Management Science*, Vol. 51/4, pp. 641-654, https://doi.org/10.1287/mnsc.1040.0343. [81]

Gal, P. (2013), "Measuring Total Factor Productivity at the Firm Level using OECD-ORBIS", *OECD Economics Department Working Papers*, No. 1049, OECD Publishing, Paris, https://doi.org/10.1787/5k46dsb25ls6-en. [68]

Gans, J. (2026), "Market Power in Artificial Intelligence", *Annual Review of Economics*, https://doi.org/10.1146/annurev-economics-051624-061832. [49]

Gómez, J. and S. Palomas (2024), "The Returns of Early Adoption of Information Technologies: Order of Adoption or Level of Adoption Advantages?", *MIS Quarterly*, Vol. 48/3, pp. 1047-1076, https://doi.org/10.25300/misq/2023/17125. [75]

Goodman-Bacon, A. (2021), "Difference-in-differences with variation in treatment timing", *Journal of Econometrics*, Vol. 225/2, pp. 254-277, https://doi.org/10.1016/j.jeconom.2021.03.014. [71]

Gutiérrez, G. and T. Philippon (2019), "Fading Stars", *AEA Papers and Proceedings*, Vol. 109, pp. 312-316, https://doi.org/10.1257/pandp.20191065. [79]

Hall, R. (1988), "The Relation between Price and Marginal Cost in U.S. Industry", *Journal of Political Economy*, Vol. 96/5, pp. 921-947, https://doi.org/10.1086/261570. [62]

Haslberger, M., J. Gingrich and J. Bhatia (2023), "No Great Equalizer: Experimental Evidence on AI in the UK Labor Market", *SSRN Electronic Journal*, https://doi.org/10.2139/ssrn.4594466. [34]

Jaumandreu, J. (2025), "Productivity, competition, and market outcomes", *Macroeconomic Dynamics*, Vol. 29, https://doi.org/10.1017/s1365100525000173. [54]

Kalemli-Ozcan, S. et al. (2015), *How to Construct Nationally Representative Firm Level Data from the Orbis Global Database: New Facts and Aggregate Implications*, National Bureau of Economic Research, Cambridge, MA, https://doi.org/10.3386/w21558. [69]

Korinek, A. and J. Vipra (2024), "Concentrating intelligence: scaling and market structure in artificial intelligence", *Economic Policy*, Vol. 40/121, pp. 225-256, https://doi.org/10.1093/epolic/eiae057. [47]

Lu, Y., G. Phillips and J. Yang (2024), *The Impact of Cloud Computing and AI on Industry Dynamics and Concentration*, National Bureau of Economic Research, Cambridge, MA, https://doi.org/10.3386/w32811. [56]

Marar, S. (2024), "Artificial Intelligence and Antitrust Law: A Primer", *SSRN Electronic Journal*, https://doi.org/10.2139/ssrn.4745321. [46]

McElheran, K. et al. (2025), *The Rise of Industrial AI in America*, Elsevier BV, https://doi.org/10.2139/ssrn.5036270. [74]

McElheran, K. et al. (2025), “The Rise of Industrial AI in America: Microfoundations of the Productivity J-curve (s)”, *Center for Economic Studies (CES) Working Paper Series*, https://www.census.gov/library/working-papers/2025/adrm/CES-WP-25-27.html. [21]

Miller, D. (2023), “An Introductory Guide to Event Study Models”, *Journal of Economic Perspectives*, Vol. 37/2, pp. 203-230, https://doi.org/10.1257/jep.37.2.203. [72]

Miller, N. (2025), “Industrial organization and The Rise of Market Power”, *International Journal of Industrial Organization*, Vol. 98, p. 103131, https://doi.org/10.1016/j.ijindorg.2024.103131. [42]

Noy, S. and W. Zhang (2023), “Experimental evidence on the productivity effects of generative artificial intelligence”, *Science*, Vol. 381/6654, pp. 187-192, https://doi.org/10.1126/science.adh2586. [32]

OECD (2025), *A Decade of OECD Competition Trends, Data and Insights*, OECD Publishing, Paris, https://doi.org/10.1787/05021ff2-en. [13]

OECD (2025), “Artificial intelligence and competitive dynamics in downstream markets”, *OECD Roundtables on Competition Policy Papers*, No. 331, OECD Publishing, Paris, https://doi.org/10.1787/ccf0624a-en. [58]

OECD (2025), “Competition in artificial intelligence infrastructure”, *OECD Roundtables on Competition Policy Papers*, No. 330, OECD Publishing, Paris, https://doi.org/10.1787/623d1874-en. [88]

Pastor-Merino, A. et al. (2026), “Does AI boost firm productivity? A web scraping and LLMs approach”, *Telecommunications Policy*, Vol. 50/2, https://doi.org/10.1016/j.telpol.2025.103138. [23]

Ropele, T. and A. Tagliabracci (2026), “The economic impact of artificial intelligence: evidence from Italian firms”, *SSRN Electronic Journal*, https://doi.org/10.2139/ssrn.6666919. [24]

Schrepel, T. and J. Potts (2025), “Measuring the openness of AI foundation models: competition and policy implications”, *Information & Communications Technology Law*, Vol. 34/3, pp. 279-304, https://doi.org/10.1080/13600834.2025.2461953. [50]

Shapiro, C. and A. Yurukoglu (2026), “Trends in Competition in the United States: What Does the Evidence Show?”, *Journal of Political Economy Microeconomics*, Vol. 4/1, pp. 316-370, https://doi.org/10.1086/733772. [43]

Varian, H. (2019), “Artificial Intelligence, Economics, and Industrial Organization”, in *The Economics of Artificial Intelligence: An Agenda*, The University of Chicago Press, Chicago. [55]

Annex A. Additional Tables and Figures

Table A1. Sample selection by sector and treatment status

	Industry (B-F)	Wholeale/retail (G)	Transport & Accom (H,I)	ICT (J)	Professional serv. (L-N)	Other services (P-S)	Total	% over total observations
All treated	2231	796	210	2015	1580	274	7106	4.63%
Drop 2018	1948	722	181	1669	1345	242	6107	3.98%
Drop always treated	1826	677	175	1455	1186	220	5539	3.61%
Drop last wave	1569	569	150	1229	995	178	4690	3.05%
Balanced pairs	1355	495	134	1086	841	145	4056	2.64%
Non-treated								
Few AI, never intensive	1089	537	209	508	1312	223	3873	2.52%
Never AI	57376	29067	16326	6149	21719	11905	142542	92.85%

Note: Further observations are dropped due to fixed effects when singletons appear.[34] Last wave firms are observations treated in 2024.
Source: Authors' elaboration on matched Orbis-Lightcast data.

[34] See Correia (2017[86]) for details.

Table A2. Selection into AI adoption I: adopters are larger, more productive, and have higher markups

	Log MFP		Markup		GO		VA	
	2018	2023	2018	2023	2018	2023	2018	2023
Adopt in 2019	11.98	12.11	1.43	1.54	100.71	142.38	34.60	50.52
Adopt in 2020	11.89	11.88	1.33	1.38	67.40	76.32	24.49	27.72
Adopt in 2021	11.65	11.63	1.31	1.33	53.27	52.95	17.52	17.65
Adopt in 2022	11.68	11.79	1.31	1.37	53.53	55.85	18.81	18.26
Adopt in 2023	11.67	11.78	1.38	1.40	46.22	52.77	15.29	16.95
Never-Adopters	10.85	10.82	1.23	1.22	15.76	14.75	4.27	4.22
	L		LP		K/L		RW	
	2018	2023	2018	2023	2018	2023	2018	2023
Adopt in 2019	300	392	189.03	233.14	54.21	70.56	113.20	119.42
Adopt in 2020	237	311	165.99	234.07	45.19	62.55	95.12	98.22
Adopt in 2021	193	215	135.80	149.53	47.70	54.06	83.69	77.85
Adopt in 2022	206	196	137.19	258.59	44.96	49.22	75.75	82.92
Adopt in 2023	175	193	128.92	150.77	37.81	49.42	79.92	86.47
Never-Adopters	70	70	85.27	90.65	36.43	39.47	51.83	50.05

Note: Monetary variables are expressed in real terms. Gross output, and value added are reported in EUR millions; labour productivity, capital per worker and average wage are reported in EUR thousands per worker. Gross output, value added, employment and capital per worker are reported as medians; log MFP, markup, labour productivity and average wage are reported as means. MFP is reported in logs.
Source: Authors' elaboration on matched Orbis-Lightcast data.

Table A3. Selection into AI adoption II: differences are statistically significant both before and after AI adoption

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	ln(MFP)	ln(markup)	ln(GO)	ln(VA)	ln(L)	ln(labour productivity)	ln(K/L)	ln(real wage)
Before AI adoption	.304***	.028***	.430***	.466***	.227***	.239***	.208***	.235***
	(.013)	(.006)	(.019)	(.017)	(.013)	(.010)	(.025)	(.008)
After AI adoption	.344***	.039***	.563***	.615***	.368***	.247***	.225***	.250***
	(.016)	(.007)	(.025)	(.023)	(.020)	(.013)	(.029)	(.011)
Fixed effects								
Firm FEs	N	N	N	N	N	N	N	N
Size, age x year FEs	Y	Y	Y	Y	Y	Y	Y	Y
Country x sector2 x year FEs	Y	Y	Y	Y	Y	Y	Y	Y
Sample	All	All	All	All	All	All	All	All
Observations	647,007	520,353	647,007	647,007	647,007	647,007	647,007	647,007
Firms	150,660	120,923	150,660	150,660	150,660	150,660	150,660	150,660
Treated firms	4,317	3,927	4,317	4,317	4,317	4,317	4,317	4,317

Notes: AI adoption is defined as the first year in which the share of AI-related vacancies over total vacancies exceeds the 75th percentile of the AI-vacancy-share distribution. Sample include both AI adopters and never adopters and estimate Equation 2. Standard errors are clustered at the firm level. Significance levels: * p<0.10, ** p<0.05, *** p<0.01.
Source: Authors' elaboration on matched Orbis-Lightcast data.

Table A4. Firms highly exposed to AI experience productivity gains

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	ln(MFP)	ln(markup)	ln(GO)	ln(VA)	ln(L)	ln(labour productivity)	ln(K/L)	ln(real wage)
After AI adoption	.043*	-.023	-.021	-.035	-.089***	.054**	.055	.025
	(.026)	(.017)	(.028)	(.036)	(.030)	(.026)	(.064)	(.020)
Fixed effects								
Firm FEs	Y	Y	Y	Y	Y	Y	Y	Y
Size, age x year FEs	Y	Y	Y	Y	Y	Y	Y	Y
Country x sector2 x year FEs	Y	Y	Y	Y	Y	Y	Y	Y
Sample	High AI exposure	High AI exposure	High AI exposure	High AI exposure	High AI exposure	High AI exposure	High AI exposure	High AI exposure
Observations	6,541	6,379	6,541	6,541	6,541	6,541	6,541	6,541
Firms	1,361	1,336	1,361	1,361	1,361	1,361	1,361	1,361
Treated firms	1,008	980	1,008	1,008	1,008	1,008	1,008	1,008

Notes: AI adoption is defined as the first year in which the share of AI-related vacancies over total vacancies exceeds the 75th percentile of the AI-vacancy-share distribution. The sample is restricted to observations in the fourth quartile of the AI-exposure distribution and includes AI-adopters only. Columns (1)-(8) are estimated using the Callaway and Sant'Anna (2021[73]) estimator and correspond to Equation 3. Standard errors are clustered at the firm level. Significance levels: * p<0.10, ** p<0.05, *** p<0.01.
Source: Authors' estimation on matched Orbis-Lightcast data.

Table A5. AI adoption disproportionately benefits larger firms

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	ln(MFP)	ln(markup)	ln(GO)	ln(VA)	ln(L)	ln(labour productivity)	ln(K/L)	ln(real wage)
baseline size: <50 employees	-.019	-.019*	.000	-.017	.008	-.025	-.024	.002
	(.023)	(.011)	(.022)	(.028)	(.025)	(.024)	(.043)	(.015)
additional effect: 50–249 employees	.009	.011	.002	.011	-.000	.011	-.001	-.003
	(.026)	(.013)	(.024)	(.032)	(.027)	(.027)	(.047)	(.017)
additional effect: 250–999 employees	.028	.015	-.011	.011	-.027	.038	.002	-.015
	(.026)	(.013)	(.024)	(.032)	(.027)	(.028)	(.048)	(.017)
additional effect: 1000+ employees	.049*	.022*	.014	.052	.002	.050	.024	-.008
	(.029)	(.013)	(.027)	(.036)	(.032)	(.032)	(.049)	(.021)
Fixed effects								
Firm FEs	Y	Y	Y	Y	Y	Y	Y	Y
Size, age x year FEs	Y	Y	Y	Y	Y	Y	Y	Y
Size category x year FEs	Y	Y	Y	Y	Y	Y	Y	Y
Country x sector2 x year FEs	Y	Y	Y	Y	Y	Y	Y	Y
Sample	AI adopters	AI adopters	AI adopters	AI adopters	AI adopters	AI adopters	AI adopters	AI adopters
Observations	26,727	24,231	26,727	26,727	26,727	26,727	26,727	26,727
Firms	5,598	5,087	5,598	5,598	5,598	5,598	5,598	5,598
Treated firms	4,263	3,875	4,263	4,263	4,263	4,263	4,263	4,263

Note: AI adoption is defined as the first year in which the share of AI-related vacancies over total vacancies exceeds the 75th percentile of the AI-vacancy-share distribution. The Table include only AI adopters, excluding firms that never adopt AI and estimate Equation 5, where Z_i^0 is baseline firm-size class. The first coefficient reports post-adoption effect for the reference category (firms with fewer than 50 employees); the interaction terms capture additional effects relative to that category. These effects are measured relative to pre-adoption firm's outcomes, conditional on firm fixed effects, country-sector-year fixed-effects, baseline-characteristic trends, including trends by baseline size category. Standard errors are clustered at the firm level. Significance levels: * p<0.10, ** p<0.05, *** p<0.01.
Source: Authors' estimation on matched Orbis-Lightcast data

Table A6. AI adoption affects sectors differently: productivity gains in ICT, markup gains in Professional Services

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	ln(MFP)	ln(markup)	ln(GO)	ln(VA)	ln(L)	ln(labour productivity)	ln(K/L)	ln(real wage)
Knowledge-intensive services								
--- ICT (J)	.043*	-.024	-.020	-.027	-.081**	.054**	-.049	.018
	(.025)	(.017)	(.031)	(.039)	(.033)	(.025)	(.062)	(.020)
--- Professional serv. (L-N)	-.022	.035*	.021	.034	.088*	-.055	-.153	-.036
	(.029)	(.019)	(.042)	(.046)	(.049)	(.036)	(.093)	(.025)
Industry (B-F)	-.033	-.002	-.063*	-.068	-.050	-.017	-.002	-.007
	(.029)	(.011)	(.035)	(.042)	(.036)	(.028)	(.046)	(.018)
Other services	-.049	-.007	-.013	-.035	-.002	-.033	.057	-.047*
	(.038)	(.013)	(.049)	(.052)	(.047)	(.035)	(.075)	(.027)
Fixed effects								
Firm FEs	Y	Y	Y	Y	Y	Y	Y	Y
Country x sector-group x year FEs	Y	Y	Y	Y	Y	Y	Y	Y
Sample	AI adopters	AI adopters	AI adopters	AI adopters	AI adopters	AI adopters	AI adopters	AI adopters
Observations	26,735	24,209	26,735	26,735	26,735	26,735	26,735	26,735
Firms	5,573	5,058	5,573	5,573	5,573	5,573	5,573	5,573
Treated firms	4045	3664	4045	4045	4045	4045	4045	4045

Note: AI adoption is defined as the first year in which the share of AI-related vacancies over total vacancies exceeds the 75th percentile of the AI-vacancy-share distribution. Coefficients are estimated using the Callaway and Sant'Anna (2021[73]) estimator and report the average post-adoption coefficient, β_{post}, from Equation 3. Each coefficient is estimated on the subsample of firms belonging to the corresponding sector. These effects are measured relative to the counterfactual evolution of firms that had not yet adopted AI, conditional on firm fixed effects, country-sector-year fixed-effects. Standard errors are clustered at the firm level. Significance levels: * p<0.10, ** p<0.05, *** p<0.01.
Source: Authors' estimation on matched Orbis-Lightcast data.

Table A7. Robustness I: definition of AI (Borgonovi et al., 2023[66]) and main countries. Signs are consistent but less statistical power

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	ln(MFP)	ln(markup)	ln(MFP)	ln(markup)	ln(MFP)	ln(markup)	ln(MFP)	ln(markup)	ln(MFP)	ln(markup)
AI adopters, before AI	.466***	.061***								
	(.018)	(.007)								
AI adopters, after AI adoption	.558***	.079***	.003	.003	.018	-.011	-.012	-.008*	-.014	-.010*
	(.021)	(.009)	(.020)	(.010)	(.035)	(.025)	(.011)	(.005)	(.011)	(.006)
After AI *x* baseline industry AI-posting Intensity *(p.p.)*							.005	-.008**		
							(.009)	(.004)		
After AI *x* baseline firm's *log* market share									.001	.002
									(.005)	(.003)
Fixed effects										
Firm FEs	N	N	Y	Y	Y	Y	Y	Y	Y	Y
Size, age x year FEs	N	N	Y	Y	Y	Y	Y	Y	Y	Y
AI-posting share x year FEs	N	N	N	N	N	N	Y	Y	N	N
Country x sector2 x year FEs	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Sample	All	All	AI adopters	AI adopters	High AI exposure	High AI exposure	AI adopters	AI adopters	AI adopters	AI adopters
Observations	440,615	352,802	13,798	12,529	3,389	3,288	14,186	12,902	11,283	10,663
Firms	108,191	86,525	2,987	2,713	733	715	3,069	2,795	2,451	2,326
Treated firms	2,981	2,650	2,029	1,831	521	505	2,228	2,019	1,758	1,654

Notes: AI adoption is defined as the first year in which the share of AI-related vacancies over total vacancies exceeds the 75th percentile of the AI-vacancy-share distribution. Columns (1) and (2) include both AI adopters and never adopters and estimate Equation 2. Columns (3)–(10) restrict the sample to AI adopters. Columns (5) and (6) further restrict the sample to observations in the fourth quartile of the AI-exposure distribution. Columns (3), (4), (5), and (6) are estimated using the Callaway and Sant'Anna (2021[73]) estimator and correspond to Equation 3. Columns (7) and (10) estimate Equation 5. Columns (7) and (8) further include trends by baseline sector-level AI-posting intensity at the two-digit sector level. In columns (7)-(8), the first coefficient report the post-adoption effect for firms in industries with average baseline AI-posting intensity, while the second coefficient shows how this effect changes when baseline industry AI-posting intensity increases by one percentage point. In columns (9)-(10), the first coefficient report the post-adoption effect for firms with average baseline log market share, while the second coefficient shows how this effect changes for a one-log-point higher baseline market share. The sample is restricted to countries with high correlation between AI-posting and AI-exposure measures and it includes: the United States, the United Kingdom, Austria, Denmark, Germany, France, Spain, Czech Republic, Hungary, Italy, Latvia, Greece and the Netherlands. Standard errors are clustered at the firm level.
Significance levels: * p<0.10, ** p<0.05, *** p<0.01.
Source: Authors' estimation on matched Orbis-Lightcast data.

Table A8. Robustness II: Countries with high correlation between the AI-posting and AI-exposure measures. Signs and statistical significance are consistent

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	ln(MFP)	ln(markup)	ln(MFP)	ln(markup)	ln(MFP)	ln(markup)	ln(MFP)	ln(markup)	ln(MFP)	ln(markup)
AI adopters, before AI	.403***	.045***	-.005	.010	-.016	.005				
	(.017)	(.007)	(.016)	(.009)	(.028)	(.022)				
AI adopters, after AI adoption	.472***	.061***	-.002	-.004	.056*	-.016	.006	-.003	-.004	-.008*
	(.022)	(.009)	(.017)	(.009)	(.031)	(.020)	(.010)	(.004)	(.010)	(.005)
After AI *x* baseline industry AI-posting Intensity *(p.p.)*							-.010	-.010***		
							(.008)	(.003)		
After AI *x* baseline firm's *log* market share									-.005	.004*
									(.005)	(.002)
Fixed effects										
Firm FEs	N	N	Y	Y	Y	Y	Y	Y	Y	Y
Size, age x year FEs	N	N	Y	Y	Y	Y	Y	Y	Y	Y
AI-posting share x year FEs	N	N	N	N	N	N	Y	Y	N	N
Country x sector2 x year FEs	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Sample	All	All	AI adopters	AI adopters	High AI exposure	High AI exposure	AI adopters	AI adopters	AI adopters	AI adopters
Observations	441,059	353,169	18,870	16,918	4,714	4,573	19,413	17,436	14,294	13,490
Firms	108,236	86,566	4,068	3,659	1,012	988	4,190	3,776	3,096	2,933
Treated firms	3,133	2,827	2,791	2,503	709	686	3,056	2,755	2,253	2,122

Notes: AI adoption is defined as the first year in which the share of AI-related vacancies over total vacancies exceeds the 75th percentile of the AI-vacancy-share distribution. Columns (1) and (2) include both AI adopters and never adopters and estimate Equation 2. Columns (3)–(10) restrict the sample to AI adopters. Columns (5) and (6) further restrict the sample to observations in the fourth quartile of the AI-exposure distribution. Columns (3), (4), (5), and (6) are estimated using the Callaway and Sant'Anna (2021[73]) estimator and correspond to Equation 3. Columns (7) and (10) estimate Equation 5. Columns (7) and (8) further include trends by baseline sector-level AI-posting intensity at the two-digit sector level. In columns (7)-(8), the first coefficient report the post-adoption effect for firms in industries with average baseline AI-posting intensity, while the second coefficient shows how this effect changes when baseline industry AI-posting intensity increases by one percentage point. In columns (9)-(10), the first coefficient report the post-adoption effect for firms with average baseline log market share, while the second coefficient shows how this effect changes for a one-log-point higher baseline market share. The sample is restricted to countries with high correlation between the AI-posting and AI-exposure measures. Countries included are: the United States, the United Kingdom, Austria, Denmark, Germany, France, Spain, Czech Republic, Hungary, Italy, Latvia, Greece and the Netherlands. Standard errors are clustered at the firm level. Significance levels: * p<0.10, ** p<0.05, *** p<0.01.

Source: Authors' estimation on matched Orbis-Lightcast data.

Table A9. Robustness III: Sectors with high correlation between the AI-posting and AI-exposure measures. Signs consistent but less statistical power

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	ln(MFP)	ln(markup)	ln(MFP)	ln(markup)	ln(MFP)	ln(markup)	ln(MFP)	ln(markup)	ln(MFP)	ln(markup)
AI adopters, before AI	.351***	.052***								
	(.016)	(.008)								
AI adopters, after AI adoption	.432***	.067***	.012	-.012	.025	-.022	.004	-.011**	-.014	-.016***
	(.020)	(.009)	(.018)	(.010)	(.028)	(.022)	(.011)	(.005)	(.009)	(.005)
After AI *x* baseline industry AI-posting Intensity *(p.p.)*							-.013	-.005		
							(.011)	(.004)		
After AI *x* baseline firm's *log* market share									-.005	.007***
									(.005)	(.002)
Fixed effects										
Firm FEs	N	N	Y	Y	Y	Y	Y	Y	Y	Y
Size, age x year FEs	N	N	Y	Y	Y	Y	Y	Y	Y	Y
AI-posting share x year FEs	N	N	N	N	N	N	Y	Y	N	N
Country x sector2 x year FEs	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Sample	All	All	AI adopters	AI adopters	High AI exposure	High AI exposure	AI adopters	AI adopters	AI adopters	AI adopters
Observations	346,155	250,133	17,597	15,704	4,102	4,009	18,080	16,169	14,552	13,706
Firms	81,351	59,107	3,697	3,312	804	793	3,808	3,420	3,060	2,893
Treated firms	3,012	2,705	2,706	2,416	627	609	2,932	2,630	2,360	2,222

Notes: AI adoption is defined as the first year in which the share of AI-related vacancies over total vacancies exceeds the 75th percentile of the AI-vacancy-share distribution. Columns (1) and (2) include both AI adopters and never adopters and estimate Equation 2. Columns (3)–(10) restrict the sample to AI adopters. Columns (5) and (6) further restrict the sample to observations in the fourth quartile of the AI-exposure distribution. Columns (3), (4), (5), and (6) are estimated using Callaway and Sant'Anna (2021[73]) estimator and correspond to Equation 3. Columns (7) and (10) estimate Equation 5. Columns (7) and (8) further include trends by baseline sector-level AI-posting intensity at the two-digit sector level. In columns (7)-(8), the first coefficient report the post-adoption effect for firms in industries with average baseline AI-posting intensity, while the second coefficient shows how this effect changes when baseline industry AI-posting intensity increases by one percentage point. In columns (9)-(10), the first coefficient report the post-adoption effect for firms with average baseline log market share, while the second coefficient shows how this effect changes for a one-log-point higher baseline market share. The sample is restricted to sectors with high correlation between the AI-posting and AI-exposure measures. Sectors included are, following NACE Rev. 2 industries: J62, J58, M, J61, C26, D, C20, C27, C21, H52, H50, B05, H51, N, E, B09, H49, C22, C25, F, I and C16. Standard errors are clustered at the firm level.
Significance levels: * p<0.10, ** p<0.05, *** p<0.01.
Source: Authors' estimation on matched Orbis-Lightcast data.

Table A10. Sectoral AI frontrunners and high-market-share firms, effect of AI on profitability

	(1)	(2)
	ln(ebitda/GO)	ln(ebitda/GO)
AI adopters, after AI adoption	-.001	-.004
(effect at the mean)	(.015)	(.016)
After AI *x* baseline industry	-.027**	
AI-posting Intensity *(p.p.)*	(.012)	
After AI *x* baseline firm's		-.001
log market share		(.006)
Fixed effects		
Firm FEs	Y	Y
Size, age x year FEs	Y	Y
AI-posting share x year FEs	Y	N
Country x sector2 x year FEs	Y	Y
Sample	AI adopters	AI adopters
Observations	23,686	18,822
Firms	5,181	4,091
Treated firms	3,759	2,991

Notes: AI adoption is defined as the first year in which the share of AI-related vacancies over total vacancies exceeds the 75th percentile of the AI-vacancy-share distribution. Columns include only AI adopters, excluding firms that never adopt AI and estimate Equation 5 that includes firm and country-sector-year fixed effects and baseline-characteristic trends. Column (1) further includes trends by baseline sector-level AI-posting intensity at the two-digit sector level. In column (1), the first coefficient reports the post-adoption effect for firms in industries with average baseline AI-posting intensity, while the second coefficient shows how this effect changes when baseline industry AI-posting intensity increases by one percentage point. In column (2), the first coefficient reports the post-adoption effect for firms with average baseline log market share, while the second coefficient shows how this effect changes for a one-log-point higher baseline market share. Standard errors are clustered at the firm level. Significance levels: * p<0.10, ** p<0.05, *** p<0.01.
Source: Authors' estimation on matched Orbis-Lightcast data.

Table A11. Heterogeneity by baseline industry AI-user and AI-developer intensity

	(1)	(2)	(3)	(4)	(5)	(6)
	ln(MFP)	ln(markup)	ln(MFP)	ln(markup)	ln(MFP)	ln(markup)
AI adopters, after AI adoption	.002	-.006*	.002	-.006	.002	-.006
(effect at the mean)	(.008)	(.004)	(.008)	(.004)	(.008)	(.004)
After AI *x* baseline industry	-.003	-.003**			-.003	.000
AI-users Intensity *(p.p.)*	(.003)	(.001)			(.005)	(.003)
After AI *x* baseline industry			-.001	-.001**	-.000	-.001
AI-developers Intensity *(p.p.)*			(.001)	(.001)	(.002)	(.001)
Fixed effects						
Firm FEs	Y	Y	Y	Y	Y	Y
Size, age x year FEs	Y	Y	Y	Y	Y	Y
Country x sector2 x year FEs	Y	Y	Y	Y	Y	Y
Sample	AI adopters	AI adopters	AI adopters	AI adopters	AI adopters	AI adopters
Observations	25,733	23,477	25,733	23,477	25,733	23,477
Firms	5,376	4,920	5,376	4,920	5,376	4,920
Treated firms	4,112	3,755	4,112	3,755	4,112	3,755

Notes: AI adoption is defined as the first year in which the share of AI-related vacancies over total vacancies exceeds the 75th percentile of the AI-vacancy-share distribution. Columns include only AI adopters, excluding firms that never adopt AI and estimate Equation 5 that includes firm and country-sector-year fixed effects, baseline-characteristic trends. In Columns (1-2), the first coefficient reports the post-adoption effect for firms in industries with average baseline AI-users intensity, while the second coefficient shows how this effect changes when baseline industry AI-users intensity increases by one percentage point. In Column (3-4), the first coefficient reports the post-adoption effect for firms with average baseline AI-developers intensity, while the second coefficient shows how this effect changes when baseline industry AI-developers intensity increases by one percentage point. In Columns (5-6), the first coefficient reports the post-adoption effect for firms with average baseline AI-users intensity and average baseline AI-developers intensity, while the other coefficients show how this effect changes when the baseline industry intensity measure increases by one percentage point while controlling for the other intensity measure. Standard errors are clustered at the firm level. Significance levels: * p<0.10, ** p<0.05, *** p<0.01.
Source: Authors' estimation on matched Orbis-Lightcast data.

Table A12. Heterogeneity by baseline firm's age.

	(1)	(2)
	ln(MFP)	ln(markup)
AI adopters, after AI adoption: 0-5 years	-.032	-.023
	(.054)	(.038)
AI adopters, after AI adoption: 6-10 years	.073	.015
	(.048)	(.030)
AI adopters, after AI adoption: 11-20 years	-.004	-.002
	(.028)	(.014)
AI adopters, after AI adoption: 21+ years	-.009	-.007
	(.020)	(.009)
Fixed effects		
Firm FEs	Y	Y
Country x sector x year FEs	Y	Y
Size x year FEs	Y	Y
Sample	AI adopters	AI adopters
Observations	26,038	23,576
Firms	5,447	4,944
Treated firms	3,938	3,566

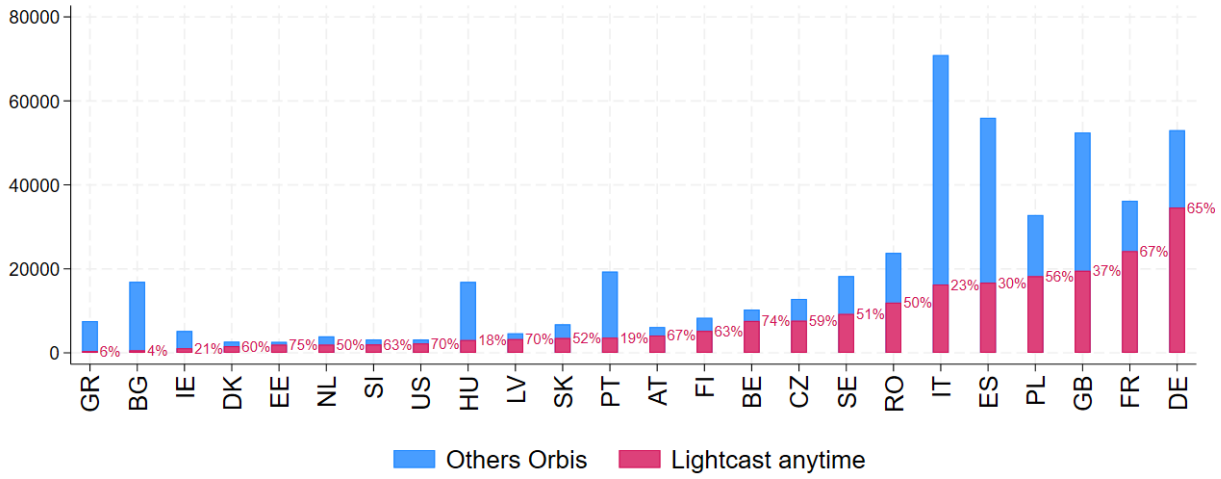
Note: AI adoption is defined as the first year in which the share of AI-related vacancies over total vacancies exceeds the 75th percentile of the AI-vacancy-share distribution. Coefficients are estimated using the Callaway and Sant'Anna (2021[73]) estimator and report the average post-adoption coefficient, β_{post}, from Equation 3. Each coefficient is estimated on the subsample of firms belonging to the corresponding age category. Baseline firm's age is measured with the first non-missing value. These effects are measured relative to the counterfactual evolution of firms that had not yet adopted AI, conditional on firm fixed effects, country-sector-year fixed-effects, and baseline size trends. Standard errors are clustered at the firm level. Significance levels: * p<0.10, ** p<0.05, *** p<0.01.
Source: Authors' estimation on matched Orbis-Lightcast data.

Table A13. No evidence of a linear relationship between peer exposure and post-adoption effects on MFP or markups.

	(1)	(2)
	ln(MFP)	ln(markup)
AI adopters, after AI adoption at 0% pre-adoption peer exposure	.005	.002
	(.012)	(.005)
AI adopters, after AI adoption × pre-adoption peer exposure *(p.p.)*	.000	-.001
	(.001)	(.001)
Fixed effects		
Firm FEs	Y	Y
Size, age x year FEs	Y	Y
Pre-adoption peer exposure × year effects	Y	Y
Cohort-specific post-adoption effects	Y	Y
Country x sector2 x year FEs	Y	Y
Sample	Adopters with measures peer expsoure	Adopters with measured peer expsoure
Observations	24,725	23,122
Firms	5,158	4,842
Treated firms	3,970	3,717

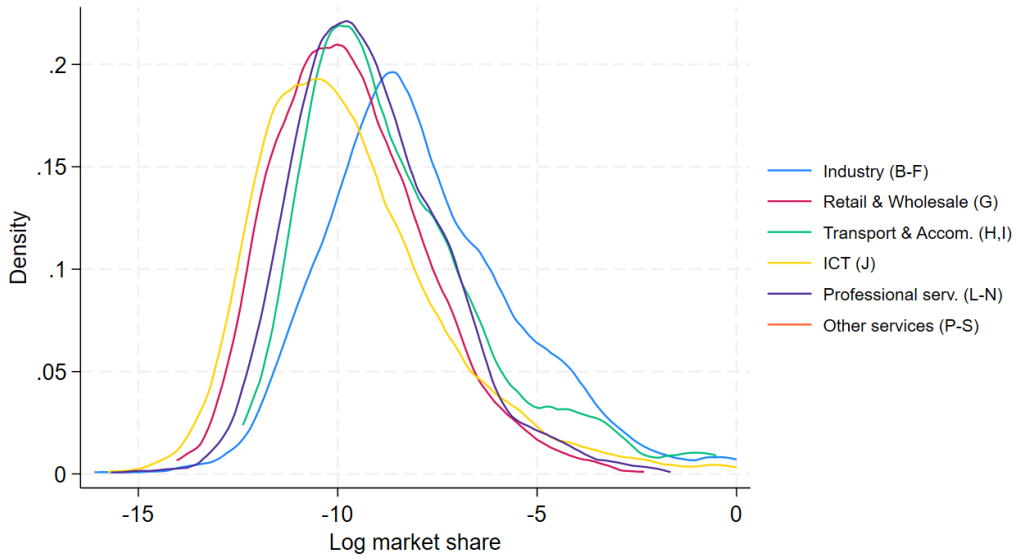
Notes: AI adoption is defined as the first year in which the share of AI-related vacancies over total vacancies exceeds the 75th percentile of the AI-vacancy-share distribution. Columns include only AI adopters for which it was possible to compute the peer exposure measure, excluding firms that never adopt AI and estimate Equation 5 that includes firm and country-sector-year fixed effects, baseline-characteristic trends, pre-adoption peer exposure trends, cohort-specific after-adoption effects. Peer exposure is the share of other firms already treated in the relevant market-year defined according to firms' geographic market scope, in the last measurable year in which the firm was still untreated . The focal firm is excluded from both the numerator and the denominator. The first coefficient reports the post-adoption effect for firms with peer exposure equal to zero, and is standardised to the treatment-cohort distribution of the estimation sample. The second coefficient shows how the post-adoption effect changes when peer exposure increases by one percentage point. Standard errors are clustered at the firm level. Significance levels: * p<0.10, ** p<0.05, *** p<0.01.
Source: Authors' estimation on matched Orbis-Lightcast data.

Figure A1. Number of firms ever matched with Lightcast in Orbis



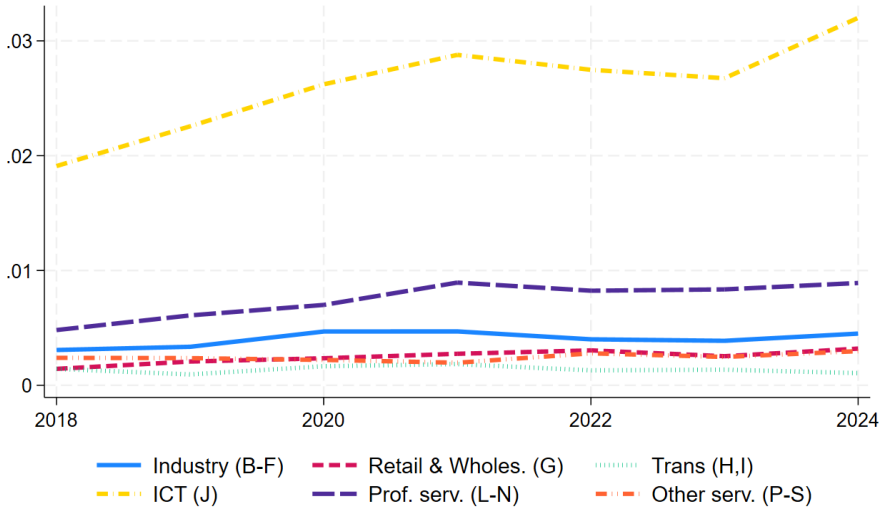
Notes: Countries are included in the sample if at least 500 firms are matched or the matching rate is at least 5%.
Source: Authors' elaboration on matched Orbis-Lightcast data.

Figure A2. Baseline firm-level market share distribution.



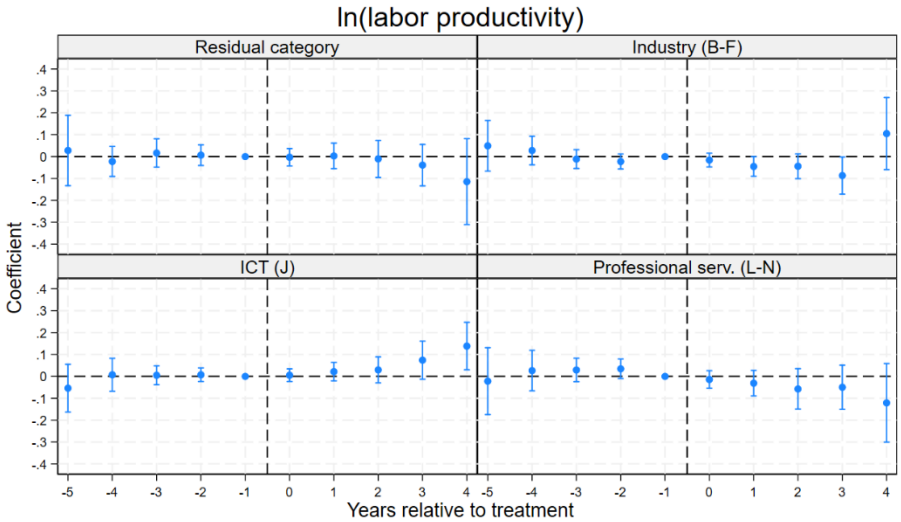
Notes: Sample includes AI adopting firms in the period 2018-2024. To construct firm-level market share, real firm-level gross output from Orbis is used as the numerator. The denominator is obtained by aggregating real annual gross output at the level of the relevant market. Relevant markets are defined at the Nace Rev. 2 two-digit level, distinguishing between domestic and internationally tradable markets.
Source: Authors' estimation on Orbis data.

Figure A3. Average firm-level share of AI-related vacancies, by sector



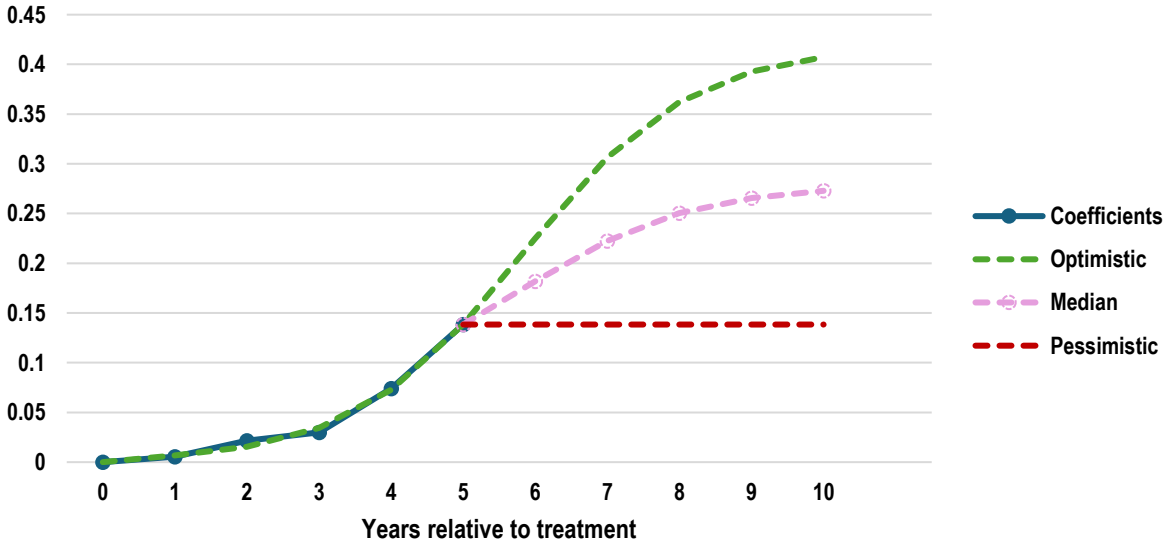
Note: Sample of firms matched in Lightcast and Orbis.
Source: Authors' elaboration on matched Orbis-Lightcast data.

Figure A4. Dynamic event-study by sector: labour productivity gains in ICT, J-Curve in Industry



Notes: AI adoption is defined as the first year in which the share of AI-related vacancies over total vacancies exceeds the 75th percentile of the AI-vacancy-share distribution. Coefficients are estimated using the Callaway and Sant'Anna (2021[73]) estimator and report the dynamic event-study effects from Equation 4. Each coefficient corresponds to a specific event time relative to the adoption year, with $(t = -1)$ omitted as the reference period. Estimates are computed separately for firms belonging to each sector group and are measured relative to the counterfactual evolution of firms that had not yet adopted AI, conditional on country-sector-year fixed-effects. Standard errors are clustered at the firm level. Confidence intervals at the 95% level.
Source: Authors' estimation on matched Orbis-Lightcast data.

Figure A5. Dynamic firm-level effects on labour productivity and projection over 10 years



Notes: the blue series reports the estimated effect of AI adoption on firms' labour productivity. AI adoption is defined as the first year in which the share of AI-related vacancies over total vacancies exceeds the 75th percentile of the AI-vacancy-share distribution. Coefficients are estimated using the Callaway and Sant'Anna (2021[73]) estimator and report the average post-adoption coefficients from equation (3). Each coefficient is estimated on the subsample of firms belonging to the corresponding sector. The pessimistic (red line) scenario assumes that no further gains occur after the last estimated dynamic coefficient. The optimistic scenario (green line) assumes gains keep rising following a sigmoid fitted on the empirically estimated coefficients in the first 5 years. The median scenario (pink line) is the median of the two scenarios. Source: Authors' estimation on matched Orbis-Lightcast data.

Annex B. Markups derivation

Deriving markups from cost minimisation

The expression for markups follows from the firm's cost minimisation problem, building on the approach of Hall (1988) as extended to the firm level by De Loecker and Warzynski (2012[61]). Consider a firm i at time t with a production technology that combines labour L_{it}, capital K_{it}, and intermediates M_{it}, to produce output Y_{it}, given firm-specific productivity Ω_{it}. Intermediates are purchased at price P_{it}^{M} and are assumed to be chosen freely each period without intertemporal considerations (as labour is subject to additional adjustment frictions). Capital is a dynamic state variable predetermined at the start of each period. The production function takes the general form:

$$Y_{it} = F(\Omega_{it}, L_{it}, K_{it}, M_{it}) \qquad \text{Equation 6}$$

Firms choose inputs optimally given their production technology, and the first-order condition for the flexible input holds. The first-order condition of this problem with respect to intermediates is:

$$P_{it}^{M} = \lambda_{it} \frac{\partial Y_{it}(.)}{\partial M_{it}}, \qquad \text{Equation 7}$$

where the Lagrange multiplier λ_{it} is the marginal value of the objective function if the output constraint is relaxed(i.e., the marginal cost of producing an additional unit of output). Multiplying both sides by the quantities P_{it} Y_{it} M_{it}, rearranging terms, and noting that the markup is defined as price over marginal cost, $\mu_{it} = P_{it}/\lambda_{it}$, yields:

$$\mu_{it} = \left(\frac{\partial Y_{it}(.)}{\partial M_{it}} \frac{M_{it}}{Y_{it}}\right) \frac{P_{it} Y_{it}}{P_{it}^{M} M_{it}} = \frac{\alpha_{M}}{S_{it}}. \qquad \text{Equation 8}$$

The markup is thus the ratio of the output elasticity of the variable input, α_{ft}^{M}, to that input's cost share in total revenue, S_{ft}. Crucially, this expression is obtained purely from the cost side and requires no assumptions on demand conditions or the nature of competition among firms.

Estimating markups

Empirically, markups are recovered from balance sheet data on revenues and input expenditures by estimating the output elasticity of intermediates, α^{M}, separately for each three-digit industry. The approach rests on two main assumptions: that at least one input can be adjusted freely within the period, and that firms are price takers for that input. Beyond these, a functional form for the production function must be chosen. Following standard practice, a Cobb-Douglas specification is adopted in each three-digit NACE Rev.2 industry:

$$Y_{it} = \Omega_{it} M_{it}^{\alpha_M} L_{it}^{\alpha_L} K_{it}^{\alpha_K}, \qquad \text{Equation 9}$$

Which in logs becomes:

$$y_{it} = \alpha_M m_{it} + \alpha_L l_{it} + \alpha_K k_{it} + \omega_{it} \quad \text{Equation 10}$$

Production function parameters are estimated using the control function approach of Ackerberg, Caves and Frazer (2015[85]), following the implementation in De Loecker and Warzynski (2012[61]). The procedure delivers industry-specific estimates of all input elasticities, including $\hat{\alpha}_M$. Given that the intermediate cost shares S_{it}^M are directly observed in the data, estimated markups are computed as:

$$\hat{\mu}_{it} = \frac{\hat{\alpha}_M}{S_{it}^M} \quad \text{Equation 11}$$

In the baseline specification, production functions are estimated separately for each three-digit industry and held constant over time, pooling all sample years. This choice is motivated on several grounds. First, theoretically, output elasticities should reflect the underlying technology of the industry rather than cyclical conditions, and given the short time span considered, it seems reasonable to assume time-invariant elasticities. Second, practical considerations play an important role: the Orbis database exhibits significant changes in sample composition over time, with coverage expanding considerably in more recent years, and minimum observation requirements for consistent production function estimation create a trade-off between industry granularity and temporal disaggregation; the baseline resolves this in favour of the finest industry granularity.

Sample construction for markup estimation

Estimating firm-level markups from production function parameters requires reliable data on revenues and input costs, specifically employee compensation, intermediates expenditure, and capital. All relevant variables are expressed in real terms using purchasing-power-parity-adjusted, country-industry-year deflators sourced from OECD STAN and OECD SNA, available at the two-digit industry level; where two-digit deflators are unavailable, more aggregated values are used. Gross output is deflated by the gross output deflator, intermediates by the intermediate input deflator, and investment and capital by the gross fixed capital formation deflator, all converted to 2015 euro prices.

The intermediates variable is not taken directly from Orbis but is constructed via the standard accounting identity (intermediates = operating revenue turnover − EBITDA − cost of employees), which draws on variables with substantially better coverage (Gal, 2013[68]). Where values for key variables such as value added or intermediates costs remain missing, they are imputed where possible using related variables available in the dataset: missing observations are replaced by the lagged value of the variable multiplied by the annual growth rate of its closest available counterpart, following (Gal, 2013[68]). Coverage in Orbis is uneven across firm sizes, with smaller firms systematically underrepresented (Bajgar et al., 2020[70]). Since reporting thresholds in many countries exclude firms with fewer than 20 employees, observations corresponding to firms whose average employment over the sample period falls below this threshold are dropped, in line with Calligaris et al. (2024[10]). The final sample for estimation includes all the countries reported in footnote 9, with the exception of Greece and Latvia, due to data limitations.

Annex C. Implementation with the Callaway and Sant'Anna estimator for event-study specification

Equation 3 and Equation 4 summarize the model underlying our analysis, i.e. an event-study specification. The average and dynamic treatment effects are estimated using the difference-in-differences estimator of Callaway and Sant'Anna (2021[73]). This estimator is suited to staggered adoption settings because it allows treatment effects to differ across adoption cohorts and over time since adoption.

Let G_i denote the first year in which firm i adopts AI. For each adoption cohort g and calendar year t, the estimator recovers group-time average treatment effects of the form:

$$ATT(g,t) = E[Y_{it}(\mathrm{g}) - Y_{it}(0) \mid G_i = g]$$

Equation 12

Under the conditional parallel trends assumption, the unobserved change in untreated potential outcomes in year t for firms first treated in year g, $E[\,Y_{it}(0) \mid G_i = g\,]$, is recovered from the observed change in outcomes of firms that have not yet adopted AI by year t.

The group-time treatment effects are then aggregated in two ways. First, the program computes an average post-adoption effect:

$$ATT^{post} = \sum_{g} \sum_{t \geq g} w_{g,t}\, ATT(g,t)$$

Equation 13

where $w_{g,t}$ are aggregation weights. In a similar way, it computes dynamic event-study effects by aggregating group-time effects, ATT_e, according to event time $e = t - g$.

This implementation avoids relying on already-treated firms as controls for later-treated firms and is robust to heterogeneous treatment effects across cohorts and over time.