Physical AI*

Adair Morse
University of California - Berkeley
& NBER

Amrita Nain
University of Iowa

June 2026

Abstract

Physical AI is the incorporation of artificial intelligence advances into the physical uses. In patent data, physical AI is as prevalent as technology AI, carries excess financial market value, and has a transatlantic prevalence pattern that may explain recent transatlantic equity return differentials. Tracing effects through real-economy inputs and production outputs, physical AI innovation raises annual employment growth by 0.72%, capital by 1.4%, and output by 0.79%, while leaving revenue TFP and cost efficiency unchanged, making the role of physical AI unique from firm automation and AI adoption. These results are robust across Jordà (2005) local projections, Sun and Abraham (2021) staggered estimators, and an IV strategy exploiting laboratory (federal and university) AI knowledge. Growth operates through both product and process innovation channels and is driven predominantly by non-automation applications of AI knowledge.

Keywords: Physical AI, Artificial Intelligence (AI), patent knowledge absorption, patent valuation, innovation dissemination, employment and AI, automation, innovation gap

* Author contacts: adair@berkeley.edu and amrita-nain@uiowa.edu. We thank Renping Li, Gerald Hoberg, and seminar participants at the University of Iowa.

I. Introduction

Since 2005, the S&P 500 has delivered an annualized return of 10.7%, compared to 7.8% for the STOXX Europe 600.[1] Much of this differential has been attributed to the United States' continuing lead in producing frontier information technology (IT).[2] We propose an additional candidate driving the current transatlantic divergence: U.S. dominance in physical AI innovation, the process by which advances in artificial intelligence (AI) are incorporated into invention in the real economy outside the IT sector. Our summary statistics document that physical AI is prevalent in patenting and unique among patents in its excess market value implications. In this paper, we study how physical AI innovation impacts the production input side (the factors of production and their productivity) and the production output side (output and cost efficiencies) of firm dynamics. Our results point to physical AI operating primarily through factors of production and output expansion rather than cost efficiency gains. This latter finding has direct implications for how the AI revolution will shape employment, growth, and cross-country equity returns.

Given the rapid advancement of AI, numerous studies examine the direct impact of firms' incorporation of AI through human capital substitution, task augmentation, and other productivity and efficiency mechanisms.[3] Simultaneously, macroeconomic studies viewing AI as a technology shock to the innovation process itself, examine the consequences of AI for aggregate growth and market value.[4] A notable few papers start the connection of these dots. For example, Babina, et al, (2024), and Fedyk, Mihet, Rishabh, and Gomes (2024), and Argente, Baslandze, Hanley, and Moreira (2025) study the role of AI on innovation generally with a focus on the employment effects, productivity, and growth. Our innovation comes with a similar focus on the innovation side of AI but our novelty is in studying the AI function as a catalytic technology for the innovation across all industry sectors of the real economy.

The scale of physical AI innovation is striking. Eighty percent of patenting occurs outside technology patent classes. We document that AI-absorbing innovation outside the IT sector is

[1] As of March 2026

[2] Brynjolfsson and Hitt (2003); Jorgenson et al. (2008); Bloom, Sadun, Van Reenen (2012); Syverson (2013) Andrews et al. (2016); Oliner and Sichel (2017); Autor et al. (2021); Draghi (2024); Huang, Papanikolaou, and Kogan (2025)

[3] Acemoglu, Autor, Hazell, and Restrepo, 2022; Eisfeldt, Schubert, and Zhang, 2023; Babina, Fedyk, He, and Hodson, 2024; Hampole, Papanikolaou, Schmidt, and Seegmiller, 2025; Brynjolfsson, Chandar, and Chen, 2025; Albanesi, Dias da Silva, Jimeno, Lamo, and Wabitsch, 2025.

[4] Acemoglu, 2003; Kogan, Papanikolaou, Seru, and Stoffman, 2017; Autor, Dorn, Katz, Patterson, and Van Reenen, 2020; Kelly, Papanikolaou, Seru, and Taddy, 2021; Acemoglu, 2025.

comparable in magnitude to AI innovation itself, implying that the economic footprint of AI extends far beyond the patents that originate it. Moreover, U.S. firms hold a 50-percentage-point larger share of AI-absorbing patents than European firms - a gap that is growing sharply, and that stands in contrast to the U.S.-Europe gaps in total innovation and in AI patenting, both of which have remained relatively stable over the same period. Using the Kogan, Papanikolaou, Seru, and Stoffman (2017) method, we show that markets price this distinction: AI-absorbing patents command a 40-basis-point premium relative to comparable patents within the same class and year. Taken together, these facts suggest that AI absorption is a key driver of cross-country growth differentials and may account for a portion of the sustained outperformance of U.S. equity markets. Yet despite its scale and apparent value, the consequences of AI-absorbing innovation for firm production, growth, and labor demand remain largely unstudied.

To fix ideas, we provide a real economy example of a physical AI invention: namely, a 2021 microsurgery device patent (US 10,895,742).[5] The device uses optical imaging to capture the surgical field and applies AI-driven real-time image processing to isolate and display magnified views of the area of surgical interest. The patent is not an AI patent but builds on earlier advances in biometrics and facial-recognition technologies – using AI advancements – that improved image identification precision (2012 patent US 8,315,440 and 2014 patent US 8,824,779). By absorbing AI technologies, the device can dynamically adjust to changing visual conditions during surgery, improving precision and adaptability across procedures and enabling remote surgical guidance and telemedical support.

The example makes two points. First, an innovation that absorbs AI as knowledge diffusion need not be classified as an AI technology; AI's influence on invention reaches well beyond its source patents. Second, the value created by AI absorption need not operate through labor substitution; it may arise instead through the expansion of what firms can invent and produce. The latter point motivates a direct comparison with Babina et al. (2024), the closest antecedent to our work.

Babina et al. (2024) show that firms investing in AI-skilled employees produce more product innovation, which drives firm growth and employment. Our measure departs from theirs in a fundamental way. Rather than capturing AI investment through hiring (an input measure), we

[5] Motivated by the attention in Goldfarb, Taska, and Teodoridis. 2020 and the volume contributions in Agrawal, Gans, Goldfarb, and Tucker, eds. 2024, we focus the example on healthcare.

measure physical AI as AI absorption through patent citations -- output evidence that AI knowledge has been incorporated into non-AI innovation. We show that the overlap between the two measures is, interestingly, quite small.[6] On one side, AI hiring may not capture the full range of innovation work especially in non-IT firms whose physical AI innovation only partially involves AI expertise.[7] Conversely, citation-based absorption may miss early-stage AI investment that has not yet materialized in patents. The two measures are likely complements, not substitutes.

To understand economic consequences, we examine how physical AI innovation alters the firm's production frontier. We trace its effects through real-economy inputs (labor scale and capital expenditure) and production outputs (revenues and profitability), while explicitly decoupling these outcomes from changes in productivity and cost efficiency. In doing so, we contribute to the growing body of research studying AI's consequences for firm growth, productivity, and asset dynamics.[8] Our study of physical AI innovation is a distinct channel addressing a large missing piece of, in the language of Babina (2026)'s survey, accurately interpreting AI's broader economic consequences.

We identify physical AI innovation in patent citations, specifically as non-IT patents which absorb AI as measure by their citations to AI. This approach follows a large literature that interprets patent citations as proxies for knowledge spillovers and technological knowledge flows (e.g., Jaffe, 1986; Jaffe, Trajtenberg, and Henderson, 1993; Hall, Jaffe, and Trajtenberg, 2001). Our analysis uses the universe of patents granted by the United States Patent and Trademark Office (USPTO) between 1979 and 2024. Using International Patent Classification (IPC) codes, we identify AI patents and measure physical AI innovation as instances in which non-IT patents cite these AI patents. We document a steady increase in physical AI patents over our sample period, with U.S. firms especially dominant in this type of innovation. In the last two years of our sample period, about 10% of non-IT patents issued to U.S. firms are AI-citing.

Our first set of empirical analyses are aimed at the factors of production (labor and capital) and factor productivity (revenue TFP). Estimating local projection regression (LPR) models of

[6] Our work is also closely tied to Fedyk, et al (2024), who study a particular instance of the AI-innovation nexus through human capital investment via data security specialists' role in non-security R&D teams.
[7] Corporate engineers (not so different from research economists) may learn AI skills sufficiently to deploy prior AI advancements in their non-AI innovation. Alternatively, akin to standard corporate practice to contract litigation lawyers rather than hire such skills in-house for one-off needs, the R&D departments may simply contract expertise to incorporate AI expertise in innovation.
[8] E.g., Kelley, Papanikolaou, Seru, and Taddy, 2021; Babina, et al, 2024; Hampole, Papanikolaou, Schmidt, and Seegmiller, 2025; Baslandze, Liu, Sojli, and Tham, 2025; Argente et al, 2025; Fedyk, et als, 2024; and Baslandze et al (2026).

Jorda (2005), we find that a one standard deviation increase in physical AI innovation results in higher factor input growth -- employment and capital -- over a five-year horizon, but with no effect on the revenue TFP. These initial results may reflect firm selection on the availability of profitable innovation opportunities and on their stock of knowledge capital, which enables them to recognize and successfully exploit the potential value of AI-using physical innovation. To address such concerns, we first estimate an interaction-weighted (IW) staggered treatment estimator of Sun and Abraham (2021). This approach improves the local projection design by estimating cohort-specific treatment effects and aggregating them using cohort sample shares. It also imposes a more structured comparison, i.e., to firms that have never engaged in innovation absorbing AI, rather than relying on the dynamic control group implicit in local projections. Consistent with the baseline results, we continue to find a significant positive effect of AI-absorbing innovation on employment and capital, but no effect on revenue TFP.

Even under this more stringent design, our estimates may reflect an endogenous firm-level decision related to firms' superior knowledge of AI-related opportunities and their ability to translate them into growth through AI-absorbing innovation. Hence, our second strategy to address endogeneity abstracts from firm decisions altogether by implementing an instrumental variable strategy that exploits localized knowledge spillovers. We instrument firms' physical AI innovation with the local stock of AI patents produced by universities and federal laboratories in the same commuting zone, building off Jaffe, Trajtenberg, and Henderson (1993) and Helper, Makan, and Shoag (2025). Conditional on industry and region controls, the local supply of AI knowledge from academic and government research institutions is plausibly unrelated to firm-specific shocks affecting outcomes, satisfying the exclusion restriction. Although most inventors are located in the same commuting zone as the assignee firm, roughly 20% of patents involve inventors located elsewhere. This creates within–commuting-zone variation in the instrument across firms, allowing us to include commuting zone times year fixed effects. In the IV specification, we find that a firm's having physical AI innovation results in 0.72% higher employment growth (as a simple 5-year annual average) and 1.4% annual capital growth over the same period.

Having traced physical AI innovation effects through real-economy inputs (labor scale and capital expenditure), while separating these outcomes from changes in productivity, we now turn to production outputs (revenues and profitability), while taking a lens to cost efficiency versus revenue-based growth. We repeat the linear projection, staggered treatment effects, and IV

estimation methodologies, finding robust evidence that physical AI innovation leads to higher growth in output and profits, but has no effect on cost efficiency. Importantly, our estimate magnitudes line up with those from the inputs into production. We find that a standard deviation more of physical AI innovation implies 0.79% higher output (revenues minus a change in inventory) and 1.15% higher profitability.

In traditional R&D theory, product innovation drives revenue expansion, while process innovation targets cost reductions (Cohen and Klepper, 1996; Klepper, 1996). However, recent work by Baslandze, Liu, Sojli, and Tham (2025) argues that process innovation can also stimulate growth by unlocking technologies and efficiencies that enable future product innovation. Motivated by these arguments, we examine whether the pro-growth effects of physical AI innovation are due to product innovation, process innovation, or both. Following Bena, Ortiz-Molina, and Simintzi (2022) and Bena and Simintzi (2025), we classify patents as product or process innovations using the text of patent claims. We find evidence that physical AI product innovation and physical AI process innovation both imply higher employment growth and output growth, robustly. What is perhaps most interesting is that the growth effects of physical AI process innovation accumulate gradually, reaching their full magnitude only over the longer horizons.[9]

Our final set of tests further confirm that the value generated by physical AI innovation is not primarily driven by labor-replacing automation. While textual analysis of patents reveals that 35% of physical AI patents relate to automation, we find that non-automation physical AI innovation strongly predicts growth in employment, capital, output, and profit. In contrast, physical AI "smart automation" shows weaker and less consistent effects. Overall, these findings suggest that the economic value of AI absorption arises primarily through non-automation channels.

Our paper contributes to the growing literature studying the economic consequences of artificial intelligence by shifting attention from the production of AI technologies to their assimilation into innovation outside the IT sector. Existing work largely studies either the labor market effects of AI adoption (e.g., Acemoglu et al., 2022; Babina et al., 2024; Hampole et

[9] These findings resonate of Babina et al. (2024), who show that hiring AI talent increases product innovation. Importantly, physical AI innovation is distinct from AI hiring: using their data, we find only a modest correlation (about 0.09) between the two measures, and our results remain unchanged when controlling for the number of AI employees. The delayed positive role for process innovation is supportive of the Baslandze et al (2025) argument that the effect of process innovation on growth is through *future* product innovation.

al., 2025) or the role of AI innovation in driving technological progress and growth. Overlapping the agenda in these studies, our paper specifically draws out the class of physical AI innovation, new to the literature.

Our paper also contributes to the literature on automation and labor markets. This literature spans early work on skill-biased technological change (Acemoglu 1998; Autor, Katz, and Krueger 1998), task-based frameworks in which automation substitutes for labor in some tasks while creating new ones (Autor, Levy, and Murnane 2003; Acemoglu and Restrepo 2019), and empirical studies using robots, AI adoption, and online vacancies (Acemoglu and Restrepo 2020; Acemoglu et al. 2022; Babina et al. 2024). While this literature focuses primarily on the adoption of automation technologies within firms or occupations, we highlight a complementary channel: the assimilation of AI knowledge into the innovation activities of non-IT firms. In a related study, Baslandze et al. (2026) use executive survey data to document that AI investment raises revenue-based productivity primarily through innovation and demand channels, with little evidence of near-term aggregate employment decline. Our work also fails to find that AI results in employment decline; however, physical AI works through an augmentation of investment into the factors of production – employment and labor – not their productivity.

More broadly, our paper contributes to the literature on absorptive capacity and knowledge spillovers (Cohen and Levinthal, 1989; Jaffe, 1986; Zahra and George, 2002; Audretsch and Belitski, 2020). We provide new evidence on how firms combine AI and domain-specific knowledge to generate innovative capabilities, and document that U.S. firms' advantage over Europe in this process represents a new mechanism through which absorptive capacity may drive cross-country differences in innovation and productivity.

The rest of the paper is organized as follows. Section II discusses the patent data and characterizes graphical punchlines of physical AI innovation. Section III presents our firm-level data. Section IV studies the results relating to real-economy inputs (labor and capital) plus these factors' productivity. Section V analyzes the flip side, the production outputs of revenue, profits, and cost efficiency. Section VI decomposes physical AI effects into concepts of process versus product patents and automation. Section VII concludes.

II. Physical AI Patent Data and Value

We begin by constructing the measure of physical AI innovation as the absorption of AI in invention, motivating its importance, and measuring its market value using the Kogan et al (2017) estimator.

II.a Patents & AI Data

We download 8.4 million utility patents granted by the U.S. Patent and Trademark Office (USPTO) between 1979 and 2024 from USPTO PatentsView.[10] The dashed line in Figure 1.A depicts the patenting (counts per grant year) over time, growing steadily to reach approximately 380,000 per year near the end of the sample. We are ultimately interested in the absorbing of AI knowledge across the breadth of the economy (i.e., in manufacturing, construction, agriculture, etc.) rather than the absorbing of AI knowledge into IT advancements themselves. Thus, in Figure 1.B, we plot non-IT patents over our sample period, which reach 240,000 per year at the sample end.[11] A first take-away from these simple graphs is the importance of the non-IT sector in overall innovation. Over time, the 6.7 million non-IT patents represent almost 80% of the total. While media attention is concentrated on innovation in the IT sector, the non-IT sector accounts for the majority of patenting activity in the U.S. economy. Therefore, understanding how these sectors absorb and translate AI knowledge into value and employment creation is of firm-order importance.

Figures 1.A and 1.B also plot patents by U.S.-based assignees (blue line) and Europe-based assignees (red dotted line). We highlight these two geographies to motivate a second contribution of our work. Since 2005, the S&P 500 has delivered an annualized return of 10.7%, compared to 7.8% for the STOXX 600 (as of May, 2026). Much of this differential has been attributed to the United States's superior innovative performance. Understanding the sources of innovation-driven value creation thus speaks not only to the U.S.–Europe performance gap but may also point toward solutions, from the European perspective. Our focus on non-IT sectors is particularly relevant here, as European firms hold comparative advantages in precisely those industries.

In Figures 1.C and 1.D, we plot the innovation gap — defined as the U.S. share of patents granted minus the European share — for all patents and non-IT patents respectively. For both total

[10] Utility patents account for 90% of USPTO patents and the term is used to distinguish from design patents. Note that in all graphs, the final year or two is more volatile from the flow of reporting and may not reflect the final counts in the period.

[11] IT patents are patents in the following IPC classes - G06F, G06Q, G06N, G06K, G06T, H04L, H04W, H04M. See appendix Table A2 for a description of these IPC codes.

and non-IT patenting, the U.S. advantage has declined since its mid-1990s peak, but the decline is concentrated in non-IT sectors. Our second graphical punchline is that Europe has been closing the gap in non-IT innovation.[12]

Turning to innovativeness in AI, in Figure 2.A and 2.B, we plot total AI patents granted[13] and total counts of AI-absorbing patents, respectively. Following the standard of measuring knowledge absorption with citations, we define AI-absorbing patents (or, interchangeably, physical AI innovation) as non-IT patents that cite a prior AI patent.[14] Our third graphical punchline lies in the similarity of the y-axis magnitudes across these two figures: patenting that absorbs AI into the physical economy is of the same order as AI patenting itself. This is unsurprising in one sense; a breakthrough technology like AI – much like the steam engine – inevitably spawns innovation across many sectors. Yet the returns to this spawning have received far less attention than the returns to AI innovation itself, both in the media and in academic work seeking to understand consequences for firm capital and workers.

Figures 2.C and 2.D turn to the U.S.–Europe gaps in AI innovation and physical AI innovation, delivering our fourth graphical punchline. The magnitude of the transatlantic AI inventiveness gaps, both approximately at 0.5, exceeds the magnitude of the gaps we plotted in Figures 1.C and 1.D. However, they diverge in trajectory. Whereas the AI innovation gap is stable, the physical AI transatlantic innovation gap is widening. We regard this divergence as an important new stylized fact. The U.S. advantage over Europe is not simply in producing AI, but in absorbing AI into the broader economy - and that absorption gap is growing.[15] To the extent that AI-absorbing

[12] One concern is that this pattern could reflect European firms increasingly filing non-IT patents at the USPTO over time rather than actual convergence in innovative activity; we address this in Appendix B using European Patent Office (EPO) data.

[13] AI patents are defined as patents with International Patent Classification (IPC) codes G06N3, G06N5, G06N7, G06N10, G06F15, and G06K9. . Note that all AI classes are nested within IT classes listed in Footnote 11.See appendix Table A2 for a description of these IPC codes.

[14] Note that we include non-technology services in our definition, because they are more than naught service processes feeding into a real economy application.

[15] AI-absorbing patents are identified using citations to U.S.-granted patents only, since USPTO PatentsView does not standardize the identifiers of cited foreign patents, precluding matching to IPC classes. This raises a home bias concern: European firms filing USPTO patents may disproportionately cite EPO-issued AI patents rather than USPTO ones, biasing their AI-absorbing patent counts downward. We address this using PATSTAT Global, which standardizes foreign patent identifiers, and find that the U.S. advantage in AI-absorbing innovation persists when all foreign citations are included.

innovation drives firm production and valuation, this pattern carries consequences for both economies' trajectories, as well as for employment.

II.b Value of Physical AI Patents

To characterize the value implied by the assimilation of AI knowledge relative to other patents, we use the patent market value (ξ) technique of Kogan et al (2017). We create two patent-level variables. The first is *PctAIcite*, calculated as the number of citations made in a patent to prior patents (henceforth, backward citations) that are AI divided by the total number of backward citations. The second is I_{AIcite}, an indicator variable that takes the value of one if at least one backward citation is made to an AI patent. Panel A of Table 2 presents the summary statistics for PctAIcite, I_{AIcite}, $\ln(\xi)$ and other patent characteristics. All variables are described in Appendix A. Because the sample is restricted to patents issued to public firms, the sample size is smaller, covering 2.5 million USPTO patents (including IT patents) between 1979 and 2024, with the only filter being the requirement of the availability of the IPC classification for both the focal patent and at least one cited patent. The mean value of I_{AIcite} indicates that about 20% of the patents cite at least one AI patent. Many of these are likely to be cases where the focal patent is an IT patent (including focal AI patents) that cites an AI patent.

Since our interest is in assimilation of AI into innovation conducted outside the IT sector, we henceforth focus on Panel B excluding focal IT patents from the sample. In non-IT patents, the mean I_{AIcite} is 0.067; just under 7% of non-IT patents cite at least one AI patent. The mean value of PctAIcite appears low at 1.2%, lower than the indicator because PctAIcite captures the yes/no on citing any AI plus the percentage of patents cited being AI. If we look within the subset of non-IT patents with I_{AIcite}=1 (i.e. at least one AI citation), the mean value of PctAIcite is 18% (not tabulated). We note that the rest of the patent characteristics such as $\ln(\xi)$, forward and backward citations, firm size, volatility are similar across Panel A and B.

We estimate how AI absorption impacts the market value of a mean patent p granted to firm i as follows:[16]

$$ln(\xi_p) = \alpha + \beta\ PctAIcite_p + \gamma_1 X_i + \gamma_2\ ln\ Backward_p + \mu_i + \rho_{kt} + \epsilon_p, \quad (1)$$

[16] In unreported tests, we use the indicator variable I_{AIcite} instead of PctAIcite and find qualitatively similar results.

where $ln\,Backward$ is the log of the number patents cited by the patent, included because assimilation of any existing knowledge, including non-AI knowledge, can be valuable. X is a vector of firm *i*-level control variables including log market capitalization (patenting may increase with size) as of the day prior to patent grant and log idiosyncratic volatility (mechanically correlated with ξ and reflective of high-growth firms having more volatile returns and higher patenting value). Finally, we include technology class times grant year fixed effects, ρ_{kt}, to force the comparison of AI assimilation to within technology-year comparables and further include μ_i, firm-fixed effects, to make the comparison within the firm's norm.[17]

As shown in Table 2, we find that AI absorbing patents are not more valuable than other patents in general (column (1)) or for IT patents (column (3)). However, reported in column (2), within the sample of non-IT patents, AI absorbing patents (i.e., physical AI citations) command higher economic value than other non-IT patents, significant at the 1% level.[18] The inference from the β coefficient in column (2) is that a physical AI patent commands a 40 basis points higher patent value ξ if the patent cites AI at the mean value (0.18) of PctAIcite when positive AI citing happens relative to a patent that does not cite any AI.[19] These results that are reminiscent of those of Zahra and George (2002) in the market valuing transformative capability defined in their work as the ability of firms to innovate by synthesizing two unconnected sets of knowledge. Combined Table 2 and Figure 2.D show that AI absorption outside the technology space is value-enhancing and that U.S. firms are increasingly global leaders in this value-enhancing innovative activity.

III. Firm-Level Data

Our analysis centers on the role of absorbing of AI innovation in firm-level production. We begin on the input side by constructing measures of firms' production factors and the factors' productivity. We then turn to measures of the outcome of production in revenue-based output and realized cost efficiency.

[17] Results are qualitatively similar if we use firm times year fixed effects and technology class fixed effects instead (see appendix Table B2)

[18] In the appendix Table B1, we find qualitatively similar results if we use I_{AIcite} as the measure of AI absorption.

[19] All else equal, the change in $ln(\xi)$ if PctAIcite changes from 0 to the average of 0.18 is 0.18*β, which equals 0.004 using β = 0.022 from column 3 of Table 2. Since the dependent variable is the natural log of ξ, this represents a 0.4% increase in ξ.

For production inputs, capital and labor are the natural factor inputs; we measure them with fixed assets and employment from Compustat. For total factor productivity (TFP), we follow Babina et al (2024) in using revenue-based TFP for comparability. Specifically, we appeal to a Cobb-Douglas technology for a firm i's revenue (y_i) production: $y_i = a_i\, k_i^{\zeta^k}\, l_i^{\zeta^l}$, where k, l, and a are our inputs of interest – capital stock, labor, and revenue TFP respectively. We isolate productivity a as the residual ϵ_{it} of the estimation:

$$log(Y_{it}) = \mu_i + \mu_t + \zeta_s^k\, log\big(k_{i,t-1}\big) + \zeta_s^l\, log(l_{it}) + \epsilon_{it}.$$

In this specification, we allow the factor intensity parameters to vary by sector s, and include firm and time t fixed effects.

For production outcomes, we analyze consequences of the use of factors in production in the business market. We measure the outcomes as output (revenues minus a change in inventories), cost efficiencies, and profits.

Our key explanatory variable is a firm-level measure of the absorbing of AI in innovation. We define *physicalAI* as the percentage of patents in non-IT firms[20] in a firm-year that cite at least one AI patent – equivalently, the firm-year average of the patent-level indicator variable I_{AIcite}. For comparison, we also show statistics on a variable $I_{physcialAI}$ equal to one for firm-years with at least one patent that cites an AI patent, and zero otherwise.[21] Our measure differs from the existing AI-employment nexus literature in three ways that are central to our contribution: we focus on innovation activity rather than AI adoption proxied through occupational exposure, we restrict attention to non-IT sectors where AI absorption is more consequential for understanding economy-wide diffusion, and we measure knowledge assimilation through citation behavior rather than AI patenting directly - capturing firms that absorb and build on AI without necessarily being AI producers.[22]

Table 3 summarizes *physicalAI* and $I_{physicalAI}$, as well as all of the production input and output variables and controls. In the full sample, 2.6% of the stock of patents in a firm-year cite at

[20] Information technology firms are defined as firms with the following three-digit SIC codes - 357, 481, 482, 484, 489, and 737.
[21] Our results are similar if we conducted analysis with $I_{physicalAI}$ rather than *physicalAI*.
[22] Other novel AI measures include Babina et al (2024), who measure human capital AI via the hiring of AI graduates and Hampole et al, (2025) and Brynjolfsson et al (2018), who extract task-level AI exposures via job description data.

least one AI patent, with an average of $I_{physicalAI}$ of 0.067, implying that just under 7% of firm-years have at least one focal patent that cites an AI patent.

IV. Physical AI and Production Factors -- Employment, Capital, and TFP

This section begins with a baseline methodology of local projection regressions. However, firms engaging in *physicalAI* may simply have more innovative opportunities. Moreover, *physicalAI* adoption may reflect a firm's accumulated knowledge capital that determines both its ability to absorb AI and its subsequent performance. To address these endogeneity concerns, we present staggered difference-in-differences (DID) estimators in Subsection IV.b that compare firms with similar innovation trajectories, following recent methodological advances (Callaway and Sant'Anna 2021; Sun and Abraham 2021; Goodman-Bacon 2021; Borusyak, Jaravel, and Spiess 2024). Finally, to rule out the possibility that the post-*physicalAI* break reflects a confounding firm-level event that simultaneously triggers both greater *physicalAI* and changes in employment and capital, in Subsection IV.c we instrument *physicalAI* using local knowledge spillovers from universities and federal labs at the commuting zone level, following Jaffe et al. (1993) and Helper, Makan, and Shoag (2025) - sources of AI knowledge that are plausibly exogenous to any individual firm.

IV.a Local Projection Analysis

Our first methodology follows Kogan et al. (2017), Kelley, Papanikolaou, Seru, and Taddy (2021) and Baslandze et al. (2025) in estimating the effect of *physicalAI* on employment, capital stock, and TFP using the Jordà (2005) Local Projection Regressions methodology. We estimate firm i's response to *physicalAI* at horizon τ according to:

$$\ln(Y_{i,t+\tau}) - \ln(Y_{it}) = \alpha_0 + \beta_\tau physicalAI_{i,t} + \gamma_1 X'_{i,t} + \gamma_2 \ln(Y_{i,t-1}) + \delta_{j,t} + \epsilon_{i,t}. \qquad (2)$$

The dependent variable is the τ-year log change in Y, where Yis alternatively the number of employees, the capital stock, or revenue TFP and $\tau = 1, \ldots, 5$ years. Our main independent variable, *physicalAI*, is the fraction of the firm's patents filed in that year that are AI-citing. We condition on the pre-period level of outcome by including lagged log employment, capital or TFP, $Y_{i,t-1}$. Identification requires isolating the *physicalAI* effect relative to an appropriate, dynamic control group. We control for control for firm size and risk dynamics by including the log value of capital stock or employment (for when that variable is not the dependent variable) and log

idiosyncratic volatility in vector X_{it}. We also include the log of one plus the number of patents a firm is granted in the year as a measure of innovative capacity. Finally, we isolate the dynamic control group with the industry-year by including SIC2 times year fixed effects δ_{jt}, double clustering standard errors by firm and year.

Table 4 presents the local projection results for employment (panel A), capital (panel B), and revenue TFP (panel C). We find a positive *physicalAI* effect on employment that remains statistically significant over all horizons. Interpreting these coefficients in economic magnitudes, we find that a standard deviation (0.132 from Table 3) higher *physicalAI* results in a 0.40% (= 0.03 x 0.132) higher employment growth one year forward or 0.23% higher employment growth per year taking the t+5 column coefficient of 0.089 and simple averaging over five years.

Panel B shows that *physicalAI* also raises the capital stock statistically significant through a four-year horizon. The magnitudes are reassuringly similar to the employment effects: a firm having *physicalAI* implies 0.16% per year higher capital stock growth on average over four years.

Finally, Panel C shows that *physicalAI* has no statistically significant effect on revenue TFP. The picture that emerges is therefore one of factor accumulation rather than factor productivity: *physicalAI* firms grow by expanding their workforce and capital stock, not by becoming more efficient at converting inputs into revenue.

IV.b Dynamic Treatment Analysis

Recent empirical advances in the analysis of two-way fixed effects (TWFE) models, including Wooldridge (2021), Goodman-Bacon (2021), Callaway and Sant'Anna (2021), Sun and Abraham (2021), and Borusyak, Jaravel, and Spiess (2024), highlight two concerns for our baseline estimates. First, with staggered treatment timing and heterogeneous treatment effects, standard TWFE estimates can be biased because later-treated units are implicitly compared to already-treated units, contaminating the control group. Second, *physicalAI* likely reflects an evolving firm strategy that precedes the patent application itself; firms may begin reallocating resources toward AI absorption before any patent is filed, potentially generating pre-trends that violate the parallel trends assumption, and more broadly the treatment anticipation assumptions of Abbring and van den Berg (2003) and Sianesi (2004), thus complicating the interpretation of our estimates.

IV.b.1 Dynamic Treatment Methodology

To address these concerns, we estimate dynamic treatment effects using the Sun and Abraham (2021) Interaction-Weighted (IW) estimator. Unlike the Jordà linear projection method that pools all treated cohorts into a single average, the IW estimator explicitly isolates each treatment cohort. By estimating cohort-specific effects and then aggregating them using cohort sample shares as weights, this method ensures that "already-treated" units are not used as controls.

We identify when a firm is first treated, i.e. when a firm first engages in physical AI innovation, as the filing year of the firm's first patent that cites an AI patent.19F[23] We denote the treatment cohort $G_{ig} \in \{g_1, g_2, \dots, g_n\}$, an indicator for a firm being in the cohort first treated at time g, and the control group $C_i \in \{0, 1\}$, the never treated. Since the *physicalAI* strategy and innovation begin a few years before the patent application, we allow for an implementation window (Callaway and Sant'Anna, 2021) of the prior two ($\delta = 2$) years before the patent granting, such that the reference period is $g - \delta - 1$.[24]

Defining the potential outcome for firm *i* at time *t* if it were treated at time g as $Y_{it}(g)$, and the counterfactual outcome for firm *i* at time *t* if it were never treated as $Y_{it}(0)$, the observed outcome is:

$$Y_{it} = \sum_{g \in G} G_{ig}\, Y_{it}(g) \;+\; C_i\, Y_{it}(0) \tag{3}$$

With this setup, our dynamic treatment effects assumptions are threefold:

1. Conditional Parallel Trends **(**CPT**):**

$$E\left[Y_{it}(0) - Y_{g-\delta-1}(0) \middle| X, G_g = 1\right] = E\left[Y_t(0) - Y_{g-\delta-1}(0) \middle| X, C = 1\right] \tag{4}$$

2. Limited Anticipation given the Implementation Window δ:

$$Y_{it} = Y_{it}(0) \;\text{ if }\; G_{ig} = 1 \text{ and } t < g - \delta \tag{5}$$

3. Treatment Effect (ATT):

$$\tau(g, t, X) = E\left[Y_t(g) - Y_t(0) \middle| X, G_g = 1\right] \neq 0 \quad \text{for } t \geq g - \delta \tag{6}$$

[23] In keeping with our norm of studying the non-IT sectors of the economy, we require that the firm not be a software or communications firm (SIC codes 357, 481, 482, 484, 489, 737).

[24] The literature on the gestation of patents (Mansfield, Schwartz, and Wagner, 1981; Mansfield, 1988; Griliches, 1990; Cohen, Nelson, and Walsh, 2000; Heeley, Matusik, and Jain, 2007) finds that the patent development period ranges from 5 months to 4 years, on average, depending on the time period and industries studied. We choose a middle range of two years.

The Sun and Abraham method involves first estimating cohort average treatment effect on the treated at each period. The Sun and Abraham (2021) Interaction-Weighted (IW) estimator is:

$$Y_{it} = \alpha_i + \lambda_t + \sum_{g \in \mathcal{G}} \sum_{\ell \neq -\delta - 1} \theta_{g\ell} \big(G_{ig} \cdot \mathbb{1}\{t - g = \ell\} \big) + \varepsilon_{it}, \tag{7}$$

where $\mathbb{1}\{t - G_{ig} = \ell\}$ is an indicator for firm *i* at time *t* being ℓ periods away (the *relative time*) from its initial treatment time g. $\alpha_i$and λ_t are firm and time fixed effects respectively. The coefficient estimator $\hat{\theta}_{gl}$ from equation (7) is a difference-in-difference estimator of the cohort average treatment effect on the treated. The Sun-Abraham IW estimator is then calculated for each relative period ℓ by taking a weighted average of $\hat{\theta}_{gl}$ using sample shares of each cohort in the relevant period ℓ as weights.

IV.b.2 Dynamic Treatment Results

Figure 3 plots the Sun and Abraham (IW) estimates for six leads and lags for employment, capital, and TFP. The shaded region represents the innovation and implementation phase, during which firms may already be reallocating resources toward the physical AI innovation ahead of the patent application, and which we therefore exclude from identification of the treatment effect. In the post-treatment period, *physicalAI* is associated with a firms having 17% change increase employment and a 21% change increase in capital stock over the horizon studied. These numbers are in log levels, rather than change in logs from the LP analysis; hence, the magnitudes are not directly comparable.[25] Panel C again shows no effect of *physicalAI* on TFP, consistent with our local projection results and with Babina et al. (2024), who find a similar null for TFP in their study of AI hiring.

IV.c Knowledge Capital Endogeneity & IV Identification through Laboratory AI

Despite the plausibility of the assumptions underlying our dynamic treatment analysis, it could be argued that the positive response observed in employment and capital are not the effect of *physicalAI* itself, but rather the consequences of firms' accumulated knowledge capital and

[25] In addition, the Table 4 calculations may underestimate an intuitive approach of considering *physicalAI* to be a zero-one concept of a firm having *physicalAI* or not. (This will more closely parallel the treatment frame in the increment of the independent variable being a 0-1.) Taking this lens, we find that firms with *physicalAI* have a 1.8% higher employment growth and 1.2% higher capital growth per year, averaged over the 5-year horizon. We prefer the more conservative interpretation in the incremental margin of *physicalAI* using the standard deviation rather than a shock form 0 to 1.

endogenous firm-level decision-making. For example, firms with deeper expertise in adjacent technologies may be both more likely to engage in *physicalAI* and be better positioned to translate it into employment and capital growth, irrespective of the AI absorption itself. In this section, we address these concerns by re-examining employment and capital growth using an instrumental variable (IV) strategy to step away from knowledge capital decisions inside the firm.

IV.c.1 Laboratories AI: IV Assumptions

We deploy the linear propagation-instrumental variable (LP-IV) methodology of Ramey (2016). Specifically, we instrument a firm's *physicalAI* with the local stock of AI patents and AI absorbing patents generated by universities and federal laboratories in the same commuting zone (CZ)**.** This strategy is motivated by Jaffe, Trajtenberg, Henderson (1993) and Helper et al (2025), who document that knowledge spillovers are highly localized geographically**,** with firms disproportionately citing patents produced nearby.

Universities and federal laboratories produce frontier scientific knowledge that is plausibly exogenous to individual firm innovation decisions, outside of the technology sector which we exclude from our sample. A higher local stock of AI patents and AI-absorbing patents from these institutions increases a firm's exposure to AI knowledge through localized spillovers, providing a source of variation in *physicalAI* that is independent of firm-specific capabilities or strategies. We use this measure as an instrument for *physicalAI*, with first-stage relevance established empirically. Our exclusion restriction requires that, conditional on region and industry controls, the historical accumulation of AI patents in nearby academic and government institutions is unrelated to contemporaneous firm-specific shocks to employment, capital, or TFP except through the firm's AI knowledge absorption.[26]

We calculate the local AI knowledge stock (local AI stock) for patent *p* filed in year *t* as the total number of AI patents or AI-citing patents granted to universities and federal laboratories located in the same commuting zone as the inventor of patent *p* over the five years from *t-5* to *t*. If a patent has multiple inventors, we use the highest value of local AI stock among its inventors. We then collapse the patent level measure local AI stock to the firm-level by taking the mean value of this across all patents in a firm-year. We name this firm-level variable *Laboratory AI*. While most

[26] A potential violation would arise if local AI researchers at universities or federal labs were directly hired by nearby firms, in which case the instrument would operate through a labor market channel rather than a knowledge spillover channel. We assess this concern directly by examining the correlation between firm-level AI hiring and *physicalAI* (see Section VI); finding it to be low.

inventors are in the same location as the assignee (i.e. the firm), for about 20% of patents, the inventor's location differs. This allows variation in the *Laboratory AI* for firms within the same commuting zone. We estimate this Linear Projection-IV (LP-IV), defined by same LPR setup of Table 4 including a vector of firm-level controls of log value of capital stock, log idiosyncratic volatility, ln(1+Patents$_{i,t-1}$), log (Emp$_{i,t-1}$) and SIC2 times year fixed effects, $\delta_{j,t}$. We additionally add $\psi_{c,t}$, commuting zone times year fixed effects.

IV.c.2 Laboratory AI: IV Results

Results from the LP-IV analysis of the production input factors are presented in Table 5, with Panels A and B studying employment and capital stock, respectively. The first-stage Kleibergen-Paap (KP) F-statistics across all column horizons and in both panels are greater than 200. The local stock of patents granted to universities and federal labs have relevance as an instrument for *physicalAI*.

In the second stage, the coefficient on *physicalAI* is positive and statistically significant in both the employment (Panel A) and capital (panel B) panels, but most robustly across medium-term horizons up to 3 years. Yet, Table 5 differs from Table 4, and from the Sun and Abraham estimates, in magnitudes. At a three-year horizon, a one standard deviation higher *physicalAI* predicts 0.72% higher annualized increase in employment growth on average and a 1.4% higher annualized capital stock growth.

V. Physical AI and Production Outcomes -- Revenues, Cost Efficiency, and Profits

Our results thus far indicate that *physicalAI* is valued by the market and leads to growth in the factors of production — employment and capital. These findings differ from prior work on AI automation, which emphasizes productivity gains from capital augmentation and labor reallocation toward tasks that are difficult to automate (Acemoglu et al. 2022; Hampole et al. 2025). In that literature, productivity improvements drive economic growth through the long-standing capital-labor productivity channel (Solow 1957; Romer 1990; Aghion and Howitt 1992) Our results suggest a complementary mechanism: *physicalAI* can also lead to economic growth through the creation of new products or services or through realized production variable cost efficiencies, all potentially making use of the new capital and new labor.

Table 6, Panel A presents local projection estimates where the dependent variables are production outcomes rather than inputs: output (revenues minus changes in inventories), cost

efficiency (COGS/Sales), and gross profits. The right-hand side specification mirrors that of Tables 4 and 5, including the lagged log dependent variable, log idiosyncratic volatility, log of one plus the lagged patent count, and SIC2 times year fixed effects. Variable definitions appear in Appendix A and summary statistics in Table 3.

If *physicalAI* drives growth in labor and capital, as our earlier results suggest, we would expect these input expansions to eventually manifest in higher revenues or profits. Separately, if *physicalAI* improves production processes, it may reduce variable input costs relative to sales, showing up as improved cost efficiency in COGS/Sales — a channel distinct from factor accumulation. We find evidence for the first channel but not the second. *physicalAI* is associated with forward growth in gross output and profits but has no robust effect on cost efficiency as measured by COGS/Sales. This pattern suggests that the output gains from *physicalAI* flow through revenue expansion rather than input cost reduction. We investigate the robustness of these findings using the Sun and Abraham IW estimator in Figure 4 and the LP-IV in the following subsection before turning to economic magnitudes and interpretation.

Figure 4 confirms the local projection findings using the Sun and Abraham IW estimator. Output (Figure 4A) and profits (Figure 4C) exhibit positive and persistent growth following *physicalAI* treatment, while cost efficiency (Figure 4B) shows no discernible treatment effect, consistent with our interpretation that *physicalAI* drives revenue expansion rather than variable cost reduction.

Turning to the LP-IV specifications in Table 6, Panel B, we continue to find strong evidence that *physicalAI* drives profit growth consistently across all projection horizons. The output results are somewhat less uniform in statistical significance, though we find evidence of positive revenue growth at the two- and five-year cumulative horizons. The economic magnitudes are consistent with those from our employment and capital results. At a five-year horizon, output and profits grow at annualized rates of 0.79% and 1.15% respectively- comparable to the annualized LP-IV estimates for employment (0.72%) and capital stock (1.4%). The consistency of these magnitudes across both input and output specifications lends additional credibility.

Taken together, the results paint a coherent picture. Firms undertaking physical AI innovation are those with revenue growth opportunities from expansion. These opportunities require the scaling up of human and physical capital.

VI. Decomposing into Process & Product Patenting and Automation

Our results suggest that the value-enhancing effects of *physicalAI* operate primarily through an output growth channel, manifesting in higher revenues, profits, and demand for capital and labor. In this section, we situate these findings within two related literatures- process innovation and automation - and discuss how our results contribute to and differ from each.

Process innovation has traditionally been viewed as a driver of lower production costs and labor-saving efficiency, mechanisms that appear less relevant in our evidence.[27] Yet, recent evidence suggests that process innovation can be a revenue-driver. In particular, Baslandze et al. (2025) show that process innovation can lead to sustained firm growth by improving product quality and by enabling the development of new products over time. This perspective connects naturally to Babina et al. (2024), who find that firms' investments in AI hiring stimulate product innovation while having little effect on process innovation. Our findings suggest that *physicalAI* unlocks growth through channels distinct from AI-hiring. Whether this growth stems from AI-absorbing product innovation or AI-absorbing process innovation is therefore an open and important question - one we take up directly in the next sub-section.

A natural nexus also exists between automation and labor-saving efficiency, again a mechanism less central to our findings.[28] Yet, a channel from automation to revenues is easily found in the literature. For example, Autor et al (2003) and Graetz and Michaels (2018) show that although computers and robotics substitute for labor in routine tasks, they complement workers in non-routine, problem-solving tasks. This complementarity between technology and skilled workers is documented extensively in the skill-biased technological change literature and provides a plausible mechanism through which even automation-oriented innovation can drive revenue expansion.[29] Whether AI-absorbing innovation operates through an automation channel - displacing routine tasks while augmenting skilled labor - or through a non-automation channel that more directly enables new products and services is therefore an empirically open question, which we examine in below.

---

[27] Bena and Simintzi (2025) find that US firms reduced process innovation after the 1999 U.S.-China bilateral agreement that increased access to cheaper labor. Bena, Ortiz-Molina, and Simintzi (2022) increase process innovation when state-level legal changes increase labor dismissal costs. Earlier work by Cohen and Klepper (1996) and Klepper (1996) views process innovation as reducing average production costs.

[28] Autor et al (2003); Acemoglu and Restrepo (2019)

[29] Autor and Dorn, 2013; Berman, Bound, and Griliches, 1994; Berman, Bound, and Machin, 1998; Acemoglu, 1998; Autor, Katz, and Krueger, 1998

VI.a Process Patenting & Automation Data & Statistics

To identify process patents, we follow the method of Bena et al. (2022) and Bena and Simintzi (2025) and identify process patents based on a textual analysis of patent claims (see Appendix C for a description of process patent classification). A patent with more than 50% process claims is classified as a process patent while the rest are classified as product patents. Using this approach, about 35% of AI-citing patents outside the IT space are process patents. To identify automation patents, we use natural language processing to classify each patent in our sample as an automation patent based on the textual description of the patent. A detailed description of the classification of a patent as an automation patent is presented in Appendix C.

In panel A of Table 7, we document the raw statistics: of all non-IT patents, whereas 6.7% are AI citing, automation and process patenting represent 13.9% and 28.5% respectively. These statistics indicate that automation and process patenting are much broader concepts than *physicalAI*.

Panels B and C look at the overlap of automation and process patents explaining the indicator variable of an AI absorbing patent. Rather than report raw correlations (largely driven by the zero no patenting firms), we instead estimate the explanatory power (partial R-Squared) of automation and process patents in explaining the variation of AI absorbing patents (in panel B) and of *physicalAI* (at the firm level, in panel C), allowing us to control for technology or industry times year fixed effects. In the high-dimensional fixed effects specifications in column (2) of Panel B, the partial R-squared of 0.007 indicates the automation indicator uniquely explains only 0.7% of the remaining 'within-firm' and 'within-technology-year' variation in AI-citing patents. Similarly, specifications in columns (4) of Panel B, the partial R-squared of 0.001 indicates the process indicator uniquely explains only 0.1% of the remaining variation in AI-citing patents. At the firm level in Panel C, column (2) and column (4), the partial R-squared indicates that after controlling for the total number of patents and including industry times year fixed effects, automation and process patents explain about 15% and 4% of the variation in *physicalAI*, respectively.[30]

[30] In columns (5) and (6) of Table 7, Panel C we regress *physicalAI* on Babina et al.'s (2024) AI-hiring measure to address the concern that our citation-based AI absorption simply reflects the output of firms' AI talent acquisition. Although the coefficient on AI hiring is positive, the partial R-squared in column (6) indicates that AI hiring explains only 1.6% of the within-industry-year variation in *physicalAI*, after controlling for total patent count. The two

VI.b Process versus Product Patenting Results

The analysis in Table 7 indicates that *physicalAI* is not simply another proxy for previously studied notions of process innovation or automation. Having established that, we now exploit the fact that *physicalAI* patents are divisible into product and process innovation. The product innovation channel is intuitive and consistent with our results since AI-absorbing product patents directly expand what firms can make and sell. Indeed, 65% of AI-absorbing non-IT patents are product patents. For the remaining 35%, the AI-absorbing process patents, we explore whether these generate pro-growth effects through a complementary process improvement channel. We run LPR using *product physicalAI*, defined as the fraction of a firm's patents that are AI-citing product patents, and *process physicalAI,* the fraction of a firm's patents that are AI-citing process patents, as the main independent variables.[31] We control for the firm's overall focus on process innovation by include the variable *process* which is the fraction of a firm's patents that are process patents of any type.

Results are presented in Table 8. In panel A, we see that *product physicalAI* and *process physicalAI* both project forward employment growth positively. Notably, for *process physicalAI*, the statistical significance strengthens in the longer time periods. Panels C and E show that *product physicalAI* projects forward output and profit growth positively and this strong positive effect is immediately visible. *Process physicalAI* also projects positive output and profit growth but the positive effect of process innovation is delayed, showing up significantly only in periods t+3 to t+5. This delay is consistent with the Baslandze et al (2025) argument that the impact of process innovation on growth works through expanding possibilities for *future* product innovation.

The positive impact of *process physicalAI* also provides a useful complement to the findings of Babina et al., who show that AI hiring is not associated with increased process innovation. While their results suggest that AI specialists primarily contribute to product-oriented innovation, our evidence allows for the interpretation that process AI innovations are implemented

---

measures are largely orthogonal, supporting our argument that AI absorption in non-IT firms is driven primarily by domain-specific employees embedding AI tools into their innovation processes rather than by newly hired AI specialists.

[31] To construct *product physicalAI*, we interact the patent-level indicator variable for AI-citing patent (I_{AIcite}) with an indicator for a product patent, and then take the mean value of the interaction per firm-year. To construct *process physicalAI*, we interact I_{AIcite} with an indicator for a process patent, and then take the mean value of the interaction per firm-year.

by domain-specific employees- such as engineers, production specialists, or technicians—who integrate AI tools into operational workflows rather than by newly hired AI specialists themselves. That is, the growth benefits of AI may arise not only from hiring AI talent, but also from the ability of existing technical personnel to embed AI technologies into firm-specific production processes.

VI.c Automation Results

Our finding that AI-absorbing innovation generates top line growth through *product physicalAI* and through *process physicalAI*, suggests that AI-absorption into the broad economy is not just about the automation of tasks. Yet smart automation that incorporates *physicalAI* may be at work, allowing us to connect recent literature on AI with the seminal literature on routine automation.[32]

We check this by analyzing whether AI-citing patents are automation oriented through a text-analysis of the patent descriptions. We find that about 35% of AI-citing patents outside the technology space meet the definition of automation patents. Although this ratio is similar to the percentage of AI-citing process patents, there is little overlap between the two types of AI-citing patents. Only about 1.5% of AI-citing patents outside the IT space meet the definition of both process patents and automation patents.

We examine the extent to which the pro-growth effects of AI-absorbing innovation act through smart automation, defined as AI-citing automation patents innovation, process innovation, or both. We run LPR regressions again using the following two measures of *physicalAI* as the main regressors: *automation physcialAI* is our measure of smart automation, defined as the fraction of a firm's patents that are AI-citing automation patents and *non-automation physicalAI* is defined the fraction of a firm's patents that are AI-citing non-automation patents. We control for the variable *automation*, which is calculated as the fraction of a firm's patents that are automation patents of any type. This control variable captures any routine, non-AI based automation patents the firm develops.

Results are presented in Table 9. We first note that *non-automation physicalAI* has a strong positive effect on growth in employment, output, and profit. The positive effects of *automation*

[32] For recent work on the effect of AI, see Acemoglu, Autor, Hazell, and Restrepo, 2022; Agrawal, Gans, and Goldfarb, 2025; Hampole, Papanikolaou, Schmidt, and Seegmiller, 2025; Brynjolfsson, Chandar, and Chen, 2025. For the long-standing literature on routine automation, see Autor, Levy and Murnane (2003), Acemoglu and Restrepo (2019), Acemoglu and Restrepo (2020) and references therein.

physicalAI, i.e. of smart automation, are mostly at the longer horizon. One can rationalize that smart innovation pays off in longer-run growth. We can see a positive effect on employment in the t+3 and t+5 horizons. The effect on output is only weakly positive over these two horizons. The effect of *automation physicalAI* on profit is weakly positive over several horizons. We see no effect of smart automation on growth in capital, TFP or COGS/Sales. Interestingly, however, the variable *automation*, which includes routine automation, has a strong positive effect on growth in employment, profit over several horizons and strong positive effect on capital over the t+1 and t+2 horizons. Routine automation matters immediately; smart automation innovation takes longer.

Overall, our results in section VI.c suggest that *physicalAI* primarily generates value through non-automation avenues such as process and product innovation. The pro-growth effects of automation are visible more strongly in routine automation than in smart automation.

VII. Conclusion

Artificial intelligence is widely viewed as a transformative technology with potentially large implications for innovation, productivity, and employment. Existing research largely focuses on the development of AI technologies themselves or on the adoption of AI within firms through hiring and task automation. In this paper, we study a complementary channel: the assimilation of AI knowledge into innovation conducted outside the information technology sector. We refer to this as AI-absorbing innovation and measure it using citations by non-IT patents to prior AI patents.

We document the large importance of non-IT sectors, where most patenting happens. Our second motivating factoid is that that U.S. firms appear to hold an increasing advantage relative to European firms in this dimension of innovation. This result could partially explain the historic difference in stock market returns between the U.S. and Europe recently. Verifying this possibility, using the patent value framework of Kogan et al. (2017), we show that non-IT patents that absorb AI knowledge are more valuable than comparable patents, consistent with the idea that combining AI with domain-specific technologies generates economically meaningful innovations.

To understand the higher value of AI-absorbing patents, we examine firm production inputs and outcomes. Across several empirical approaches—including local projection regressions, dynamic DID estimators, and an instrumental-variable strategy based on localized AI knowledge spillovers—we find that firms engaging in AI-absorbing innovation experience higher employment and capital growth as factors of production, but not higher TFP use of these inputs.

This suggests a growth channel for physical AI rather than a productivity or cost efficiency channel. We verify this interpretation in firm outcomes: AI-absorbing innovation predicts higher growth in output and profits but not cost efficiency. These findings suggest that the economic value of AI-absorbing innovation arises primarily through revenue expansion.

Finally, we show that both product-oriented and process-oriented physical AI innovations contribute to these outcomes. Product physical AI innovations generate immediate growth effects, while process physical AI innovations produce delayed but economically meaningful gains. Overall, our results highlight the importance of understanding how general-purpose technologies such as artificial intelligence diffuse across technological domains and shape firm growth and employment.

References

Abbring, Jaap H., and Gerard J. van den Berg. 2003. “The Nonparametric Identification of Treatment Effects in Duration Models.” *Econometrica* 71 (5): 1491–517.

Acemoglu, Daron. 1998. “Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality*.” *The Quarterly Journal of Economics* 113 (4): 1055–89.

Acemoglu, Daron. 2003. “Labor- and Capital-Augmenting Technical Change.” *Journal of the European Economic Association* 1 (1): 1–37.

Acemoglu, Daron. 2025. “Nobel Lecture: Institutions, Technology, and Prosperity.” *American Economic Review* 115 (6): 1709–48.

Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo. 2022. “Artificial Intelligence and Jobs: Evidence from Online Vacancies.” *Journal of Labor Economics* 40 (S1): S293–340.

Acemoglu, Daron, and Pascual Restrepo. 2019. “Automation and New Tasks: How Technology Displaces and Reinstates Labor.” *Journal of Economic Perspectives* 33 (2): 3–30.

Acemoglu, Daron, and Pascual Restrepo. 2020. “Robots and Jobs: Evidence from US Labor Markets.” *Journal of Political Economy* 128 (6): 2188–244.

Aghion, Philippe, and Peter Howitt. 1992. “A Model of Growth Through Creative Destruction.” *Econometrica* 60 (2): 323–51.

Agrawal, Ajay, Joshua Gans, Avi Goldfarb, and Catherine E. Tucker, eds. 2024. *The Economics of Artificial Intelligence: Health Care Challenges*. National Bureau of Economic Research Conference Report. University of Chicago Press.

Agrawal, Ajay K., Joshua S. Gans, and Avi Goldfarb. 2025. “Genius on Demand: The Value of Transformative Artificial Intelligence.” Working Paper No. 34316. Working Paper Series. National Bureau of Economic Research, October.

Albanesi, Stefania, António Dias da Silva, Juan F. Jimeno, Ana Lamo, and Alena Wabitsch. 2025. “New Technologies and Jobs in Europe*.” *Economic Policy* 40 (121): 71–139.

Andrews, Dan, Chiara Criscuolo, and Peter N. Gal. 2016. “The Global Productivity Slowdown, Technology Divergence and Public Policy: A Firm Level Perspective.” *(Working Paper / OECD Background; Also Published in OECD/Brookings Volumes and Widely Cited in Top-Journal Literature)*.

Argente, David, Salomé Baslandze, Douglas Hanley, and Sara Moreira. 2025. “Patents to Products: Product Innovation and Firm Dynamics.” Working Paper No. 34592. Working Paper Series. National Bureau of Economic Research, December.

Audretsch, David B., and Maksim Belitski. 2020. “The Role of R&D and Knowledge Spillovers in Innovation and Productivity.” *European Economic Review* 123 (April): 103391.

Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen. 2020. “The Fall of the Labor Share and the Rise of Superstar Firms*.” *The Quarterly Journal of Economics* 135 (2): 645–709.

Autor, David H., and David Dorn. 2013. “The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market.” *American Economic Review* 103 (5): 1553–97.

Autor, David H., David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen. 2021. “Concentrating on the Fall of the Labor Share and Productivity: The Role of Technology and Market Power.” *Quarterly Journal of Economics* 136 (1): 1–60.

Autor, David H., Lawrence F. Katz, and Alan B. Krueger. 1998. “Computing Inequality: Have Computers Changed the Labor Market?*.” *The Quarterly Journal of Economics* 113 (4): 1169–213.

Autor, David H., Frank Levy, and Richard J. Murnane. 2003. “The Skill Content of Recent Technological Change: An Empirical Exploration.” *The Quarterly Journal of Economics* 118 (4): 1279–333.

Babina, Tania. 2026. “Understanding Firms’ AI Efforts and Their Economic Impact.” SSRN Scholarly Paper No. 6623645. Social Science Research Network, April 21.

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson. 2024. “Artificial Intelligence, Firm Growth, and Product Innovation.” *Journal of Financial Economics* 151 (January): 103745.

Bacchiocchi, Emanuele, and Fabio Montobbio. 2010. "International Knowledge Diffusion and Home-Bias Effect: Do USPTO and EPO Patent Citations Tell the Same Story?" Scandinavian Journal of Economics 112 (3): 441–470.

Baslandze, Salomé, Zachary Edwards, John R. Graham, et al. 2026. “Artificial Intelligence, Productivity, and the Workforce: Evidence from Corporate Executives.” SSRN Scholarly Paper No. 6456970. Social Science Research Network, March 1.

Baslandze, Salomé, Leo Liu, Elvira Sojli, and Wing Wah Tham. 2025. “Foundational Processes and Growth.” *FRB Atlanta Working Paper*, FRB Atlanta Working Paper, February 19, 2025–01.

Bena, Jan, Hernán Ortiz-Molina, and Elena Simintzi. 2022. “Shielding Firm Value: Employment Protection and Process Innovation.” *Journal of Financial Economics* 146 (2): 637–64.

Bena, Jan, and Elena Simintzi. 2025. “Machines Could Not Compete with Chinese Labor: Evidence from US Firms’ Innovation.” *Review of Finance* 29 (6): 1619–61.

Berman, Eli, John Bound, and Zvi Griliches. 1994. “Changes in the Demand for Skilled Labor within U.S. Manufacturing: Evidence from the Annual Survey of Manufacturers.” *The Quarterly Journal of Economics* 109 (2): 367–97.

Berman, Eli, John Bound, and Stephen Machin. 1998. “Implications of Skill-Biased Technological Change: International Evidence*.” *The Quarterly Journal of Economics* 113 (4): 1245–79.

Bloom, Nicholas, Raffaella Sadun, and John Van Reenen. 2012. “Americans Do IT Better: US Multinationals and the Productivity Miracle.” *American Economic Review* 102 (1): 167–201.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess. 2024. “Revisiting Event-Study Designs: Robust and Efficient Estimation.” *The Review of Economic Studies* 91 (6): 3253–85.

Bresnahan, Timothy F., and M. Trajtenberg. 1995. “General Purpose Technologies ‘Engines of Growth’?” *Journal of Econometrics* 65 (1): 83–108.

Brynjolfsson, Erik, and Lorin M. Hitt. 2003. “Computing Productivity: Firm-Level Evidence.” *Review of Economics and Statistics* 85 (4): 793–808.

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond. 2025. “Generative AI at Work*.” *The Quarterly Journal of Economics* 140 (2): 889–942.

Callaway, Brantly, and Pedro H. C. Sant’Anna. 2021. “Difference-in-Differences with Multiple Time Periods.” *Journal of Econometrics*, Themed Issue: Treatment Effect 1, vol. 225 (2): 200–230.

Cohen, Wesley M., and Steven Klepper. 1996. “Firm Size and the Nature of Innovation within Industries: The Case of Process and Product R&D.” *The Review of Economics and Statistics* 78 (2): 232–43.

Cohen, Wesley M., and Daniel A. Levinthal. 1990. “Absorptive Capacity: A New Perspective on Learning and Innovation.” *Administrative Science Quarterly* 35 (1): 128–52.

Cohen, Wesley M., Richard Nelson, and John P. Walsh. 2000. *Protecting Their Intellectual Assets: Appropriability Conditions and Why U.S. Manufacturing Firms Patent (or Not)*. Working Paper No. 7552. NBER Working Papers. National Bureau of Economic Research, Inc.

Criscuolo, Paola. 2006. “The ‘home Advantage’ Effect and Patent Families. A Comparison of OECD Triadic Patents, the USPTO and the EPO.” Scientometrics 66 (1): 23–41.

Eisfeldt, Andrea L., Gregor Schubert, and Miao Ben Zhang. 2023. “Generative AI and Firm Values.” Working Paper No. 31222. Working Paper Series. National Bureau of Economic Research, May. https://doi.org/10.3386/w31222.

Fedyk, Anastassia, Roxana Mihet, Kumar Rishabh, and Orlando Gomes. 2024. “Data Innovation Complementarity and Firm Growth.” SSRN Scholarly Paper No. 4559921. Social Science Research Network, September 1.

Goldfarb, Avi, Bledi Taska, and Florenta Teodoridis. 2020. “Artificial Intelligence in Health Care? Evidence from Online Job Postings.” *AEA Papers and Proceedings* 110 (May): 400–404.

Goodman-Bacon, Andrew. 2021. “Difference-in-Differences with Variation in Treatment Timing.” *Journal of Econometrics*, Themed Issue: Treatment Effect 1, vol. 225 (2): 254–77.

Graetz, Georg, and Guy Michaels. 2018. “Robots at Work.” *The Review of Economics and Statistics* 100 (5): 753–68.

Griliches, Zvi. 1990. “Patent Statistics as Economic Indicators: A Survey.” *Journal of Economic Literature* 28 (4): 1661–707.

Hall, Bronwyn H., Adam Jaffe, and Manuel Trajtenberg. 2005. “Market Value and Patent Citations.” *The RAND Journal of Economics* 36 (1): 16–38.

Hampole, Menaka, Dimitris Papanikolaou, Lawrence D. W. Schmidt, and Bryan Seegmiller. 2025. “Artificial Intelligence and the Labor Market.” Working Paper No. 33509. Working Paper Series. National Bureau of Economic Research, February.

Heeley, Michael B., Sharon F. Matusik, and Neelam Jain. 2007. “Innovation, Appropriability, and the Underpricing of Initial Public Offerings.” *The Academy of Management Journal* 50 (1): 209–25.

Helper, Susan, Resem Makan, and Daniel Shoag. 2025. *Federally-Funded Science as Engines of Regional Innovation and Prosperity: The Case of the National Labs*. Working Paper.

Huang, Qiushi, Dimitris Papanikolaou, and Leonid Kogan. 2025. “Tech Dollars: Technological Innovation and Exchange Rates.” SSRN Scholarly Paper No. 4667334. Social Science Research Network, January 1.

Jaffe, Adam B. 1986. “Technological Opportunity and Spillovers of R & D: Evidence from Firms’ Patents, Profits, and Market Value.” *The American Economic Review* 76 (5): 984–1001.

Jaffe, Adam B., Manuel Trajtenberg, and Rebecca Henderson. 1993. “Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations.” *The Quarterly Journal of Economics* 108 (3): 577–98.

Jorgenson, Dale W., Mun S. Ho, and Kiichi Stiroh. 2008. “A Retrospective Look at the U.S. Productivity Growth Resurgence.” *Journal of Economic Perspectives* 22 (1): 3–24.

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy. 2021. “Measuring Technological Innovation over the Long Run.” *American Economic Review: Insights* 3 (3): 303–20.

Klepper, Steven. 1996. “Entry, Exit, Growth, and Innovation over the Product Life Cycle.” *The American Economic Review* 86 (3): 562–83.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. 2017. “Technological Innovation, Resource Allocation, and Growth*.” *The Quarterly Journal of Economics* 132 (2): 665–712.

Mansfield, Edwin. 1988. “The Speed and Cost of Industrial Innovation in Japan and the United States: External vs. Internal Technology.” *Management Science* 34 (10): 1157–68.

Mansfield, Edwin, Mark Schwartz, and Samuel Wagner. 1981. “Imitation Costs and Patents: An Empirical Study.” *The Economic Journal* 91 (364): 907–18.

Oliner, Stephen D., and Daniel E. Sichel. 2017. “Information Technology and Productivity: Where Are We Now and Where Are We Going?” *Journal of Economic Perspectives* 31 (1): 3–30.

Ramey, Valerie. 2016. “Macroeconomic Shocks and Their Propagation.” In *Handbook of Macroeconomics*, vol. 2. Elsevier.

Romer, Paul M. 1990. “Endogenous Technological Change.” *Journal of Political Economy* 98 (5): S71–102.

Sianesi, Barbara. 2004. “An Evaluation of the Swedish System of Active Labor Market Programs in the 1990s.” *The Review of Economics and Statistics* 86 (1): 133–55.

Solow, Robert M. 1957. “Technical Change and the Aggregate Production Function.” *The Review of Economics and Statistics* 39 (3): 312–20.

Sun, Liyang, and Sarah Abraham. 2021. “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects.” *Journal of Econometrics*, Themed Issue: Treatment Effect 1, vol. 225 (2): 175–99.

Syverson, Chad. 2013. “Will History Repeat Itself? Comments on ‘Is the Information Technology Revolution Over?’” *International Productivity Monitor* 26: 6–16.

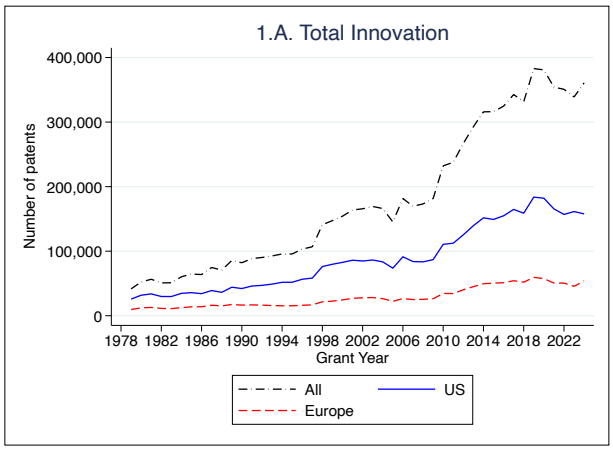
“The Draghi Report on EU Competitiveness.” n.d. Accessed May 27, 2026. https://commission.europa.eu/topics/competitiveness/draghi-report_en.

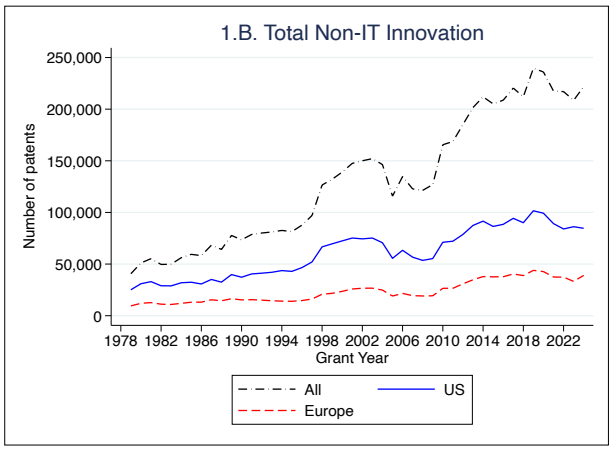
Wooldridge, Jeffrey. 2021. *Two-Way Fixed Effects, the Two-Way Mundlak Regression, and Difference-in-Differences Estimators*.

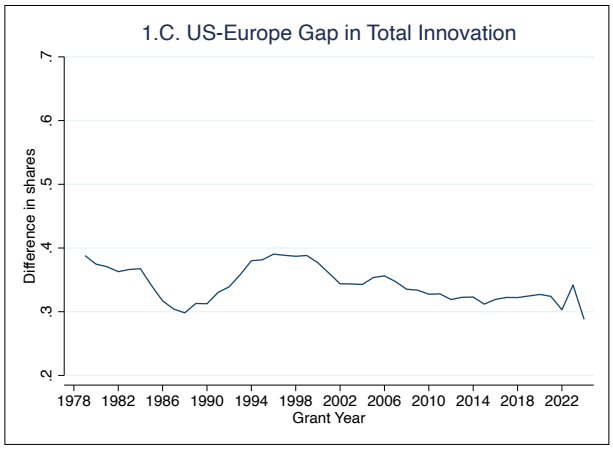
Zahra, Shaker A., and Gerard George. 2002. “Absorptive Capacity: A Review, Reconceptualization, and Extension.” *The Academy of Management Review* 27 (2): 185–203.

Figure 1: Total Innovation and non-IT Innovation

Figure 1.A. depicts the total number USPTO patents over the period 1979 to 2024 and the number of patents granted to U.S firms and European firms respectively. Figure 1.C. depicts the difference in the share of total patents held by U.S. and European companies, i.e., the difference between the blue (solid line) share of all patents and the red (dashed) line share of all patents. Figure 2.A. depicts the total number non-IT patents granted by USPTO over the period 1979 to 2024 and the number of non-IT patents granted to U.S firms and European firms respectively. Non-IT patents are patents *not* in the following information technology IPC classes: G06F, G06Q, G06N, G06K, G06T, H04L, H04W, H04M. Figure 2.C. depicts the difference in the share of not-IT patents held by U.S. and European companies, i.e., the difference between the blue (solid line) share of all patents and the red (dashed) line share of all patents.







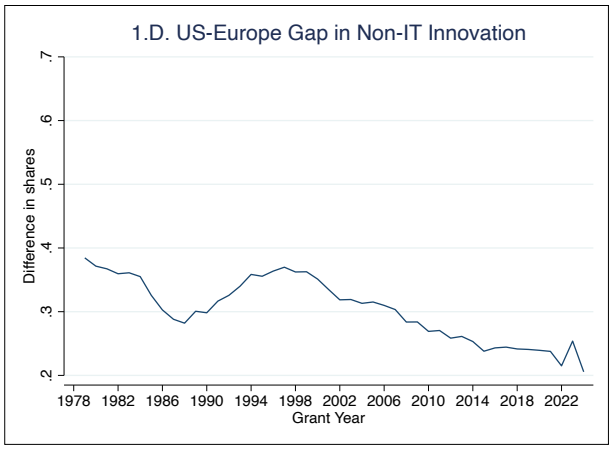
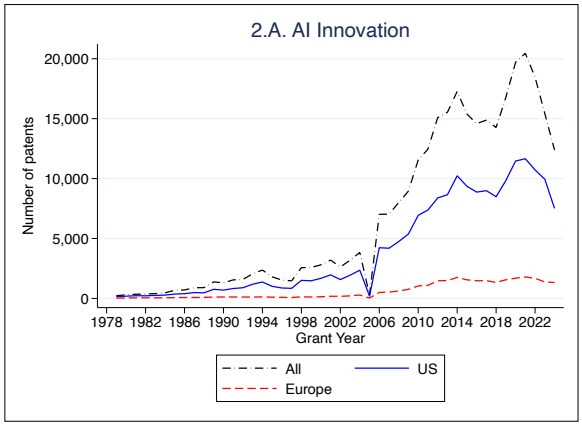
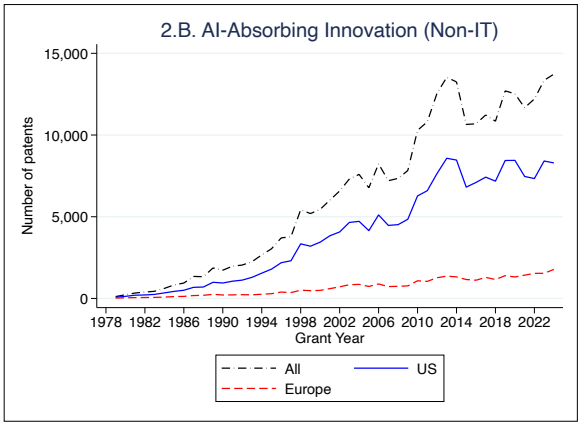
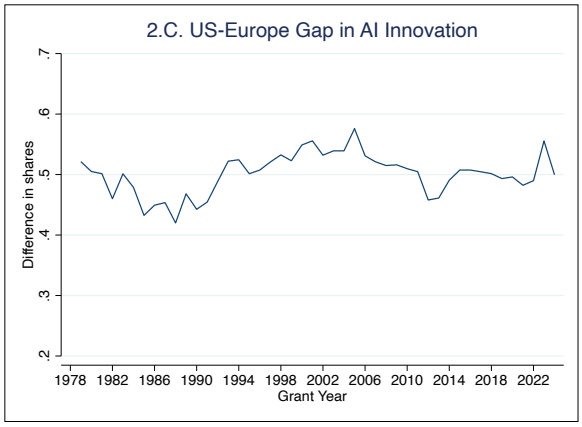


Figure 2: AI Innovation and AI-Absorbing Innovation

Figure 2.A. depicts the number of Artificial Intelligence (AI) patents granted by USPTO, and the number of AI patents granted to U.S. and European firms respectively. Patents with IPC class G06N3, G06N5, G06N7, G06N10, G06F15, and G06K9 are classified as AI patents. Figure 2.C. depicts the difference in the share of total AI patents held by U.S. and European companies. Figure 2.B. depicts AI-absorbing innovation, defined as non-IT patents that cite AI patents. Figure 1.D. depicts the difference in the share of AI-absorbing innovation by U.S. and European companies.







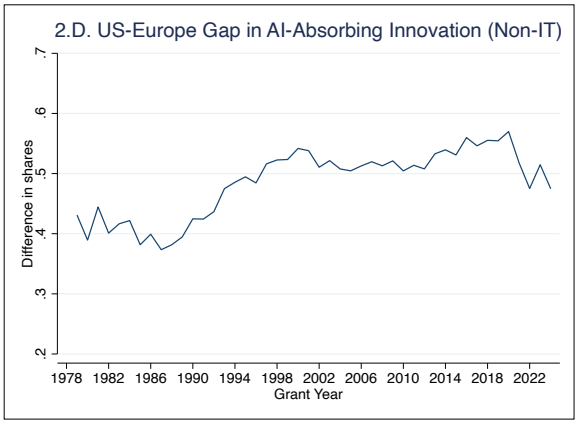


Figure 3
Two-Way Fixed Effects Analysis: Employment

This figure depicts the Sun and Abraham (2021) Interaction-Weighted (IW) estimator for the six periods before and six periods after the innovation phase. First treatment is defined as the filing year of the first AI-absorbing (AIAI) patent by a firm and the innovation phase. Using G to denote the year when a unit first becomes treated, a cohort average treatment effect on the treated, $\hat{\theta}_{g,l}$, is first estimated using the following regression

$$Y_{it} = \alpha_i + \lambda_t + \sum_{g} \sum_{\ell \neq -3} \theta_{g\ell}\left(G_{ig} \cdot \mathbb{1}\{t - g = \ell\}\right) + \varepsilon_{it}$$

Y is log employment, log capital or revenue TFP and $G_{ig} = \mathbb{1}(t - g = \ell)$ is an indicator for unit i being ℓ periods away from initial treatment at calendar time t. α_i are firm-fixed effects and λ_t are year-fixed effects. The figures show, for each relative period ℓ, the weighted average of $\hat{\theta}_{gl}$ using sample shares of each cohort in the relevant period ℓ as weights. The shaded region represents the period of anticipatory behavior during which AI-absorbing innovation is already underway.

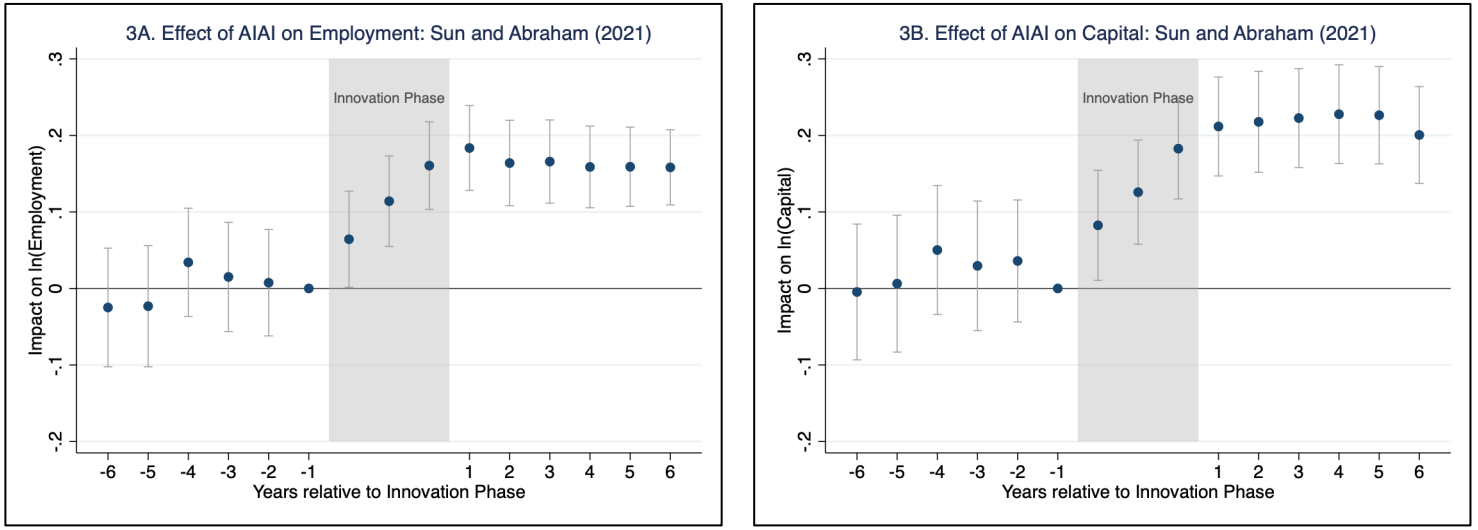


Figure 3 Continued

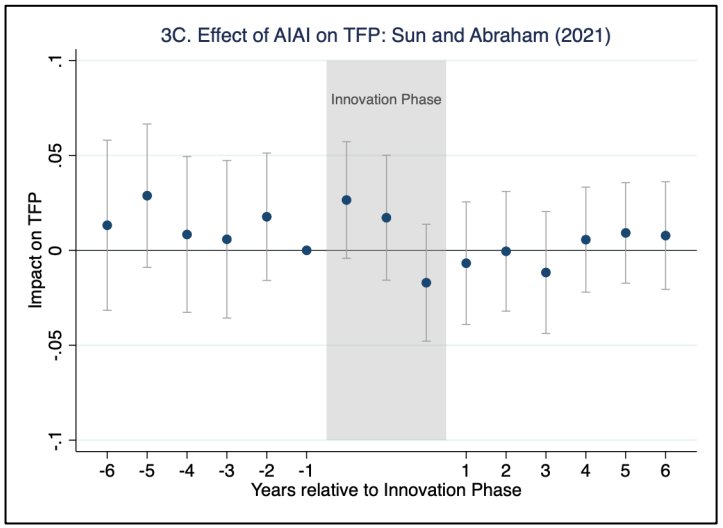
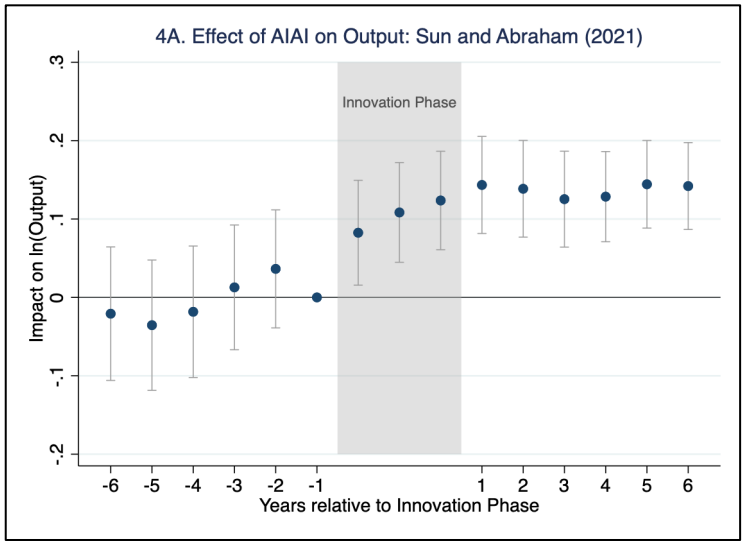


Figure 4
Two-Way Fixed Effects Analysis: Output and Profit

This figure depicts the Sun and Abraham (2021) Interaction-Weighted (IW) estimator for the six periods before and six periods after the innovation phase. First treatment is defined as the filing year of the first AI-absorbing (AIAI) patent by a firm and the innovation phase. Using G to denote the year when a unit first becomes treated, a cohort average treatment effect on the treated, $\hat{\theta}_{g,l}$, is first estimated using the following regression

$$Y_{it} = \alpha_i + \lambda_t + \sum_{g} \sum_{\ell \neq -3} \theta_{g\ell}\left(G_{ig} \cdot \mathbb{1}\{t - g = \ell\}\right) + \varepsilon_{it}$$

Y is log output, cost efficiency (COGS/Sales) or log gross profit and $G_{ig} = \mathbb{1}(t - g = \ell)$ is an indicator for unit *i* being ℓ periods away from initial treatment at calendar time *t*. α_i are firm-fixed effects and λ_t are year-fixed effects. The figures show, for each relative period ℓ, the weighted average of $\hat{\theta}_{gl}$ using sample shares of each cohort in the relevant period ℓ as weights. The shaded region represents the period of anticipatory behavior during which AI-absorbing innovation is already underway.



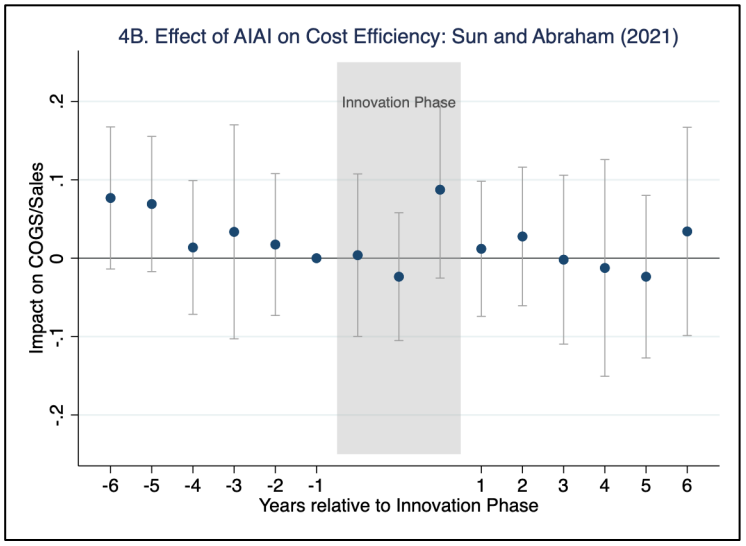


Figure 4 Continued

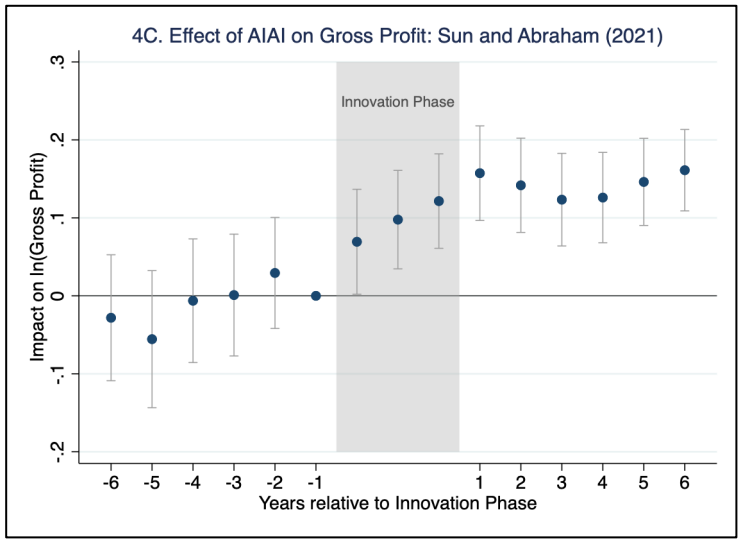


Table 1: Patent-level Summary Statistics

Panel A presents summary statistics for all utility patents issued to publicly traded firms in the U.S between 1979 and 2024. Panel B presents summary statistics for the subsample of non-IT patents only. Non-IT patents are patents *not* in the following information technology IPC classes: G06F, G06Q, G06N, G06K, G06T, H04L, H04W, H04M. $\ln(\xi)$ is the Kogan et al (2017) patent value. *Pct AI cite* is the percentage of cited patents that are AI patents. I_{AIcite} is an indicator variable taking the value 1 if the focal patent cites at least one AI patent. *Forward* captures truncation bias adjusted forward citations which involves dividing the number of forward citations a patent receives by the mean number of forward citations received by patents granted in the same year. *ln Backward* is the log of the number patents cited by the focal patent. *AI Focal Patent* is an indicator variable equal to 1 if focal patent is an AI patent. Patents with IPC class G06N3, G06N5, G06N7, G06N10, G06F15, and G06K9 are classified as AI patents. *ln Market cap* is calculated as of one day prior to patent grant date as the natural log of shares outstanding times share price. $ln(\sigma)$ is the natural log of idiosyncratic volatility in the patent grant year.

Panel A: All Patents

Variable	N	Mean	SD	p10	p25	p50	p75	p90
$\ln(\xi)$	2,533,690	0.977	2.263	-2.449	-0.168	1.445	2.509	3.463
Pct AI cite	2,533,690	0.056	0.158	0	0	0	0	0.2
I_{AIcite}	2,533,690	0.197	0.398	0	0	0	0	1
ln (1+Forward)	2,533,690	0.415	0.545	0	0	0.240	0.573	1.077
ln Backward	2,533,690	1.988	1.162	0.693	1.099	1.946	2.639	3.401
AI Focal Patent	2,533,595	0.054	0.227	0	0	0	0	0
ln Market Cap	2,345,326	16.106	2.522	12.672	14.449	16.266	18.202	19.071
$\ln(\sigma)$	2,343,675	-4.087	0.467	-4.629	-4.416	-4.136	-3.811	-3.479

Panel B: Non-IT Patents

Variable	N	Mean	SD	p10	p25	p50	p75	p90
$\ln(\xi)$	1,778,737	0.813	2.290	-2.684	-0.529	1.351	2.391	3.302
Pct AI cite	1,778,737	0.012	0.067	0	0	0	0	0
I_{AIcite}	1,778,737	0.067	0.250	0	0	0	0	0
ln (1+Forward)	1,778,737	0.413	0.528	0	0.056	0.249	0.570	1.035
ln Backward	1,778,737	2.031	1.133	0.693	1.386	1.946	2.639	3.401
AI Focal Patent	1,778,737	0	0	0	0	0	0	0
ln Market Cap	1,633,865	15.631	2.446	12.290	14.024	15.796	17.541	18.647
$\ln(\sigma)$	1,632,719	-4.063	0.473	-4.622	-4.394	-4.113	-3.792	-3.450

Table 2: Physical AI Innovation and Patent Value

This table presents results from estimating the following patent-level equation:

$$ln(\xi_p) = \alpha + \beta\ PctAIcite_p + \gamma_1 X_i + \gamma_2 \ln Backward_p + \mu_i + \rho_{k,t} + \epsilon_p$$

where ξ is the Kogan et al (2017) value of patent *p*, *PctAIcite* is the share of a patent's backward citations that cite an AI patent, *ln Backward* is the log of the number patents cited by the patent, X is a vector of firm-level control variables including log market capitalization and log idiosyncratic volatility, $\mu_i$are firm-fixed effects, $\rho_{k,t}$ are technology class times grant year fixed effects. Standard errors are clustered by year. ***, **,* indicate significance at the 1%, 5%, and 10% respectively. All variable definitions are provided in Appendix A. The sample period is 1979 to 2024.

Dependent Variable:		$ln(\xi)$	
Sample:	All Patents	Non-IT Patents	IT Patents
	(1)	(2)	(3)
PctAIcite	-0.005	0.022***	-0.003
	(-1.037)	(2.803)	(-0.657)
ln Market Cap	0.893***	0.906***	0.844***
	(75.487)	(71.692)	(57.806)
$\ln(\sigma)$	0.755***	0.761***	0.723***
	(29.152)	(26.265)	(19.610)
ln Backward	0.005***	0.007***	0.001
	(4.400)	(4.652)	(0.519)
Observations	2,336,260	1,625,472	708,417
R-squared	0.937	0.939	0.932
Firm FE	Yes	Yes	Yes
PatentClass*Year FE	Yes	Yes	Yes
Clustered SE	Year	Year	Year

Table 3: Firm-level Summary Statistics

Summary statistics for firms excluding the software and communications sector over the sample period 1979 to 2024. Physical AI Innovation is the percentage of patents in a firm-year that cite an AI patent. $I_{PhysicalAI}$ is an indicator variable equal to one for firm-years with at least one patent that cites an AI patent, and zero otherwise. ln Employment is the natural log of the number of employees. ln Capital Stock is the natural log of capital defined as Compustat *ppent*, deflated by the NIPA price of equipment. Revenue TFP is the residual from regressing log real sales on log employment and log lagged capital, controlling for firm fixed effects and year fixed effects. The regression is estimated using OLS separately for each 2-digit SIC industry. ln Output is the natural log of output defined as Compustat *sale* plus change in inventories (*invt*) deflated by the CPI. Cost efficiency is calculated as Compustat *cogs*/(*sale*+change in *invt*). ln Profit is the natural log of gross profit defined as Compustat *sale* minus *cogs*. ln (1+Patents) is the natural log of one plus the number patents filed by the firm. $\ln(\sigma)$ is log of Idiosyncratic volatility in the patent filing year

Variable	N	Mean	SD	p10	p25	p50	p75	p90
Physical AI Innovation	264,962	0.026	0.132	0	0	0	0	0
$I_{PhysicalAI}$	264,962	0.066	0.249	0	0	0	0	0
ln Employment	215,887	-0.167	2.340	-3.194	-1.814	-0.124	1.504	2.879
ln Capital Stock	220,498	3.750	2.805	0.235	1.813	3.649	5.707	7.510
Revenue TFP	145,869	0.002	0.419	-0.325	-0.129	0.008	0.148	0.345
ln Output	203,349	4.654	2.442	1.650	3.096	4.723	6.312	7.744
Cost efficiency	204,334	0.982	3.605	0.280	0.473	0.652	0.793	0.917
ln Profit	212,161	3.672	2.240	0.890	2.183	3.668	5.180	6.542
ln (1+Patents)	264,962	0.426	1.024	0	0	0	0	1.609
$\ln(\sigma)$	250,656	-3.615	0.625	-4.409	-4.082	-3.649	-3.183	-2.778

Table 4: Effect of Physical AI Innovation on Employment, Capital and Revenue TFP

This table presents the relation between physical AI innovation (denoted *PhysicalAI*) and firm-level outcomes. We present the coefficients β_τ for cumulative forward time periods $\tau \in \{1, \ldots, 5\}$ from estimations of the following model:

$$\ln(Y_{i,t+\tau}) - \ln(Y_{i,t}) = \alpha_0 + \beta_\tau PhysicalAI_{it} + \gamma_1 X'_{it} + \delta_{jt} + \epsilon_{it}$$

where Y is one of the following: the log number of employees (Panel A); log capital stock (Panel B); Revenue TFP (Panel C). X is a vector of firm-level control variables that varies by panel. Log idiosyncratic volatility, log number of patents, and lagged dependent variable are included as control variables in all panels. In Panel A, we include log of capital stock at time t as an additional control variable. In Panel B, we include log number of employees at time t as an additional control variable. In Panel C, both log of capital stock and log number of employees at time t are include as control variables. $\delta_{j,t}$ are SIC2 times year fixed effects. Standard errors are double-clustered by firm and year. ***, **,* indicate significance at the 1%, 5%, and 10% respectively. All variable definitions are provided in Appendix A. The sample period is 1979 to 2024.

	Horizon				
	t+1	t+2	t+3	t+4	t+5
	Panel A: Change in Log Employment				
Physical AI Innovation	0.030***	0.048***	0.077***	0.079***	0.089***
	(4.80)	(3.97)	(4.77)	(3.46)	(3.32)
ln (1+Patents) (lag)	0.006***	0.012***	0.018***	0.023***	0.027***
	(5.87)	(6.46)	(6.58)	(6.75)	(6.65)
ln Employment (lag)	-0.037***	-0.062***	-0.082***	-0.100***	-0.116***
	(-18.09)	(-20.22)	(-19.39)	(-17.82)	(-16.70)
Observations	164385	148238	134035	121597	110643
R-squared	0.057	0.073	0.083	0.092	0.101
	Panel B: Change in Log Capital				
Physical AI Innovation	0.025***	0.035**	0.059**	0.057*	0.055
	(3.04)	(2.20)	(2.61)	(1.91)	(1.57)
ln (1+Patents) (lag)	0.010***	0.018***	0.023***	0.029***	0.036***
	(5.67)	(6.28)	(5.92)	(6.14)	(6.21)
ln Capital (lag)	-0.061***	-0.099***	-0.128***	-0.151***	-0.171***
	(-13.46)	(-13.66)	(-14.03)	(-14.46)	(-15.54)
Observations	167201	150968	136597	123964	112783
R-squared	0.078	0.099	0.115	0.124	0.132
Industry x Year FE & Controls	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year

Table 4 continued

	Panel C: Change in Revenue TFP				
Physical AI Innovation	-0.004	0.005	0.002	0.004	-0.008
	(-0.61)	(0.47)	(0.17)	(0.35)	(-0.70)
ln (1+Patents) (lag)	0.000	0.002	0.005***	0.008***	0.010***
	(0.29)	(1.29)	(2.72)	(3.86)	(4.13)
Revenue TFP (lag)	-0.235***	-0.377***	-0.481***	-0.578***	-0.644***
	(-23.31)	(-32.19)	(-30.19)	(-38.12)	(-39.71)
Observations	114202	102183	92439	83891	76343
R-squared	0.075	0.121	0.162	0.207	0.235
Industry x Year FE & Controls	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year

Table 5: Effect of Instrumented Physical AI Innovation on Employment and Capital

This table presents an instrumental variable analysis of the relation between Physical AI Innovation and firm-level outcomes. We present the coefficients β_τ for cumulative forward time periods $\tau \in \{1, \ldots, 5\}$ from estimations of the following model:

$$\ln(Y_{i,t+\tau}) - \ln(Y_{i,t}) = \alpha_0 + \beta_\tau \widehat{PhysicalAI}_{i,t} + \gamma_1 X'_{i,t} + + \delta_{jt} + \psi_{c,t} + \epsilon_{i,t}$$

where Y is one of the following: the log number of employees (Panel A); log capital stock (Panel B). $\widehat{PhysicalAI}$ is the fitted value from a first stage regression of Physical AI Innovation on *Laboratory AI*, the number of AI patents or AI-citing patents granted to universities and federal labs located in the firm's commuting zone. X is a vector of firm-level control variables that varies by panel. Log idiosyncratic volatility, log number of patents, and lagged dependent variable are included as control variables in all panels. In Panel A (Panel B), we include log of capital stock (number of employees) at time t as an additional control variable. $\delta_{j,t}$ are SIC2 times year fixed effects and $\psi_{c,t}$ are commuting zone times year fixed effects. Standard errors are double-clustered by firm and year. ***, **,* indicate significance at the 1%, 5%, and 10% respectively. All variable definitions are provided in Appendix A. The sample period is 1979 to 2024.

	t+1	t+2	t+3	t+4	t+5
	Panel A: Change in Log Employment				
First Stage					
Laboratory AI	0.028***	0.028***	0.029***	0.029***	0.030***
First stage F-statistic	260.43	251.60	232.38	207.94	202.50
Second Stage					
$\widehat{PhysicalAI}$	0.144***	0.149***	0.163*	0.182	0.219
	(4.13)	(2.71)	(1.76)	(1.66)	(1.68)
ln Employment (lag)	-0.039***	-0.065***	-0.086***	-0.106***	-0.123***
	(-17.23)	(-19.49)	(-18.34)	(-16.94)	(-15.61)
ln (1+Patents) (lag)	0.003**	0.010***	0.016***	0.022***	0.026***
	(2.54)	(4.51)	(4.58)	(4.86)	(4.82)
Observations	140592	127019	114951	104297	94863
	Panel B: Change in Log Capital				
First Stage					
Laboratory AI	0.028***	0.028***	0.029***	0.029***	0.030***
First stage F-statistic	260.71	251.04	226.58	205.14	204.14
Second Stage					
$\widehat{PhysicalAI}$	0.249***	0.351***	0.318**	0.320*	0.236
	(4.93)	(3.90)	(2.43)	(1.83)	(1.25)
ln Capital (lag)	-0.065***	-0.105***	-0.134***	-0.156***	-0.176***
	(-14.98)	(-15.11)	(-15.29)	(-15.60)	(-16.58)
ln (1+Patents) (lag)	0.004*	0.010**	0.016***	0.022***	0.031***
	(1.74)	(2.49)	(3.10)	(3.24)	(3.86)
Observations	142514	128851	116674	105894	96296
Industry x Year FE & Controls	Yes	Yes	Yes	Yes	Yes
Commuting Zone x Year FE	Yes	Yes	Yes	Yes	Yes

Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year

Table 6: The Effect of Physical AI Innovation on Firm Output, Cost Efficiency, and Profit

This table presents the relation between physical AI innovation and firm-level output and profit. In Panel A, we present the coefficients β_τ for cumulative forward time periods $\tau \in \{1, \dots, 5\}$ from estimations of the following:

$$\ln(Y_{i,t+\tau}) - \ln(Y_{it}) = \alpha_0 + \beta_\tau \text{Physical}AI_{it} + \gamma_1 X'_{it} + \delta_{jt} + \epsilon_{it}$$

where Y are firm outcomes represented in the Panel names; *Physical*AI is the fraction of the firm's patents that are AI-citing. In Panel B, we present estimates from the following instrumental variable regression:

$$\ln(Y_{i,t+5}) - \ln(Y_{it}) = \alpha_0 + \beta_5 \widehat{\text{Physical}}AI_{it} + \gamma_1 X'_{it} + \delta_{jt} + \epsilon_{it}$$

where Y are firm outcomes represented in the Panel names; $\widehat{Physical}$AI is the fitted value from a first stage regression of physical AI innovation on the number of AI patents or AI-citing patents granted to universities and federal labs located in the firm's commuting zone. X is a vector of firm-time level control variables including log value of capital stock, the log number of employees, log idiosyncratic volatility, ln(1+Patents$_{i,t-1}$), and ln(Y$_{i,t-1}$); δ_{jt} are SIC2 times year fixed effects. Standard errors are double-clustered by firm and year. ***, **, * indicate significance at the 1%, 5%, and 10% respectively. All variable definitions are provided in Appendix A. The sample period is 1979 to 2024.

	t+1	t+2	t+3	t+4	t+5
			Panel A: LP regressions		
			A.1: Output		
Physical AI Innovation	0.036***	0.066***	0.090***	0.110***	0.122***
	(6.45)	(4.18)	(4.25)	(4.26)	(3.80)
Observations	146492	132119	120057	109327	99755
R-squared	0.075	0.101	0.107	0.114	0.122
			A.2: Cost Efficiency		
Physical AI Innovation	-0.016	0.032	0.010	-0.118**	-0.046
	(-0.68)	(0.45)	(0.13)	(-2.45)	(-0.75)
Observations	147888	133328	121099	110217	100535
R-squared	0.019	0.031	0.028	0.045	0.039
			A.3: Profit		
Physical AI Innovation	0.049***	0.083***	0.110***	0.125***	0.142***
	(6.08)	(5.38)	(5.99)	(5.37)	(5.12)
Observations	151001	136828	124365	113339	103522
R-squared	0.090	0.110	0.117	0.122	0.133
Industry x Year FE & Controls	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year

Panel B: LP-IV regressions

	B.1: Output				
$\widehat{PhysicalAI}$	0.102	0.200**	0.155	0.168	0.301**
	(1.47)	(2.45)	(1.40)	(1.25)	(2.15)
Observations	125126	112976	102722	93541	85260
First stage F-statistic	262.46	250.26	232.46	211.37	208.88
	B.2: Cost Efficiency				
$\widehat{PhysicaAI}$	0.032	0.321	-0.168	0.312	0.356
	(0.06)	(0.64)	(-0.25)	(0.54)	(0.53)
Observations	126303	113979	103580	94273	85902
First stage F-statistic	265.91	254.43	231.47	209.19	205.41
	B.3: Profit				
$\widehat{PhysicaAI}$	0.190***	0.297***	0.271***	0.339***	0.436***
	(6.07)	(4.73)	(3.38)	(3.75)	(4.37)
Observations	128675	116743	106165	96759	88301
First stage F-statistic	292.19	281.71	260.85	233.46	232.16
Industry x Year FE & Controls	Yes	Yes	Yes	Yes	Yes
Commuting Zone x Year FE	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year

Table 7: Statistics tying Automation, Process Patents and AI Hiring to Physical AI

Presented are statistics of patent-level (restricted to non-IT patent sample) and firm-level variables (restricted to non-IT firms only) for 1979-2024. I_{AIcite} is an indicator variable taking the value 1 if the focal patent cites at least one AI patent; $I_{Automation}$ is an indicator variable taking the value 1 if the focal patent is an automation patent; $I_{Process}$ is an indicator variable that takes the value 1 if the more than 50% of the focal patent's claims are process claims. Firm-level variables are: *Physical AI Innovation* is fraction of the firm's patents that are AI-citing. *Automation* is the fraction of a firm's patents that are automation patents. *Process* is the fraction of a firm's patents with more than 50% process claims. *ln AI-hiring* is the natural log of Babina et al. (2024) measure of AI employment. See Appendix C for the classification of process claims, automation patents, and non-IT patents. Panels B and C present correlational estimates to assess partial R-squared among patent taxonomy at the patent and firm level respectively. ***, **,* indicate significance at the 1%, 5%, and 10% respectively. Variable definitions are provided in Appendix A.

Panel A: Summary statistics

Variable	N	Mean	SD	p10	p25	p50	p75	p90
Patent level								
I_{AIcite}	1,778,737	0.067	0.250	0	0	0	0	0
$I_{Automation}$	1,777,127	0.139	0.346	0	0	0	0	1
$I_{Process}$	1,778,724	0.285	0.452	0	0	0	1	1
Firm level								
Physical AI Innov.	264,962	0.026	0.132	0	0	0	0	0
Automation	264,962	0.045	0.169	0	0	0	0	0.053
Process	264,962	0.072	0.213	0	0	0	0	0.281
ln AI-Hiring	40,548	0.713	1.285	0	0	0	0.881	2.492

Panel B: Patent-Level Distinction among Physical AI Innovation, Automation, and Process Patents

	(1)	(2)	(3)	(4)
	Dependent Variable: Indicator for AI Cite (I_{AIcite})			
$I_{Automation}$	0.118***	0.058***		
	(36.37)	(38.99)		
$I_{Process}$			0.024***	0.020***
			(11.68)	(26.25)
Observations	1,777,221	1,769,666	1,778,819	1,771,262
Partial R-sq	0.027	0.007	0.002	0.001
Fixed effects	No	Firm, Class x Year	No	Firm, Class x Year
Clustering	Year	Year	Year	Year

Panel C: Firm-Level Distinction among AI Absorbing, Automation, and Process Patents

	(1)	(2)	(3)	(4)	(5)	(6)
	Dependent Variable: Physical AI Innovation					
Automation	0.304***	0.283***				
	(0.013)	(0.013)				
Process			0.133***	0.128***		
			(0.011)	(0.010)		
ln AI-Hiring					0.020***	0.019***
					(0.003)	(0.003)
Observations	238197	233528	238197	233528	39,347	39,300
Partial R-sq	0.155	0.145	0.043	0.043	0.022	0.016
Fixed effects	No	SIC2 x Year	No	SIC2 x Year	No	SIC2 x Year
Control for # patents	Yes	Yes	Yes	Yes	Yes	Yes
Clustering	Firm, year	Firm, year	Firm, year	Firm, year	Firm, year	Firm, year

Table 8: The Role of Physical AI Innovation via Automation and Routine Automation

This table presents the relation between physical AI innovation, via automation and non-automation, and firm-level outcomes. We present the coefficients β_τ for cumulative forward time periods $\tau \in \{1, \dots, 5\}$ from estimations of the following model:

$$\ln(Y_{i,t+\tau}) - \ln(Y_{it}) = \alpha_0 + \beta_{1\tau} AutoPhysicalAI_{it} + \beta_{2\tau} NonAutoPhysicalAI_{it} + \beta_{3\tau} Auto_{it} + \gamma_1 X'_{it} + \delta_{jt} + \epsilon_{it}$$

where Y are firm outcomes represented in the Panel names; A$utoPhysicalAI$ is the fraction of the firm's patents that are AI-citing automation patents; $NonAutoPhysicalAI$ is the fraction of the firm's patents that are AI-citing non-automation patents. Auto is the fraction of a firm's patents that are automation patents. The classification of automation patents is described in Appendix C. X is a vector of firm-time level control variables including log value of capital stock, the log number of employees, log idiosyncratic volatility, ln(1+Patents$_{i,t-1}$), and ln(Y$_{i,t-1}$); δ_{jt} are SIC2 times year fixed effects. Standard errors are double-clustered by firm and year. ***, **,* indicate significance at the 1%, 5%, and 10% respectively. All variable definitions are provided in Appendix A. The sample period is 1979 to 2024. Panel A is the full sample of publicly traded firms. Panel B restricts firms serving as comparable to innovative firm-years defined as firm-years with at least one patent in the period *t* through *t-5*

	t+1	t+2	t+3	t+4	t+5
			Panel A: Employment		
Automation Physical AI	0.005	0.025	0.064**	0.053	0.086**
	(0.53)	(1.56)	(2.46)	(1.57)	(2.27)
Non-Automation Physical AI	0.030***	0.037**	0.053**	0.062**	0.056
	(3.56)	(2.32)	(2.56)	(2.04)	(1.54)
Automation	0.023***	0.033***	0.034**	0.042**	0.033
	(3.73)	(3.09)	(2.38)	(2.27)	(1.55)
Observations	164385	148238	134035	121597	110643
R-squared	0.057	0.073	0.083	0.092	0.101
			Panel B: Capital		
Automation Physical AI	0.010	-0.017	0.038	0.033	0.080
	(0.57)	(-0.63)	(1.00)	(0.68)	(1.48)
Non-Automation Physical AI	0.020*	0.038*	0.046	0.052	0.020
	(1.84)	(1.89)	(1.65)	(1.45)	(0.45)
Automation	0.020*	0.049***	0.033	0.028	0.007
	(1.97)	(3.03)	(1.44)	(1.01)	(0.21)
Observations	167201	150968	136597	123964	112783
R-squared	0.078	0.099	0.115	0.124	0.132
			Panel C: Output		
Automation Physical AI	0.009	0.031	0.069**	0.074*	0.102*
	(0.63)	(1.19)	(2.55)	(1.91)	(1.99)
Non-Automation Physical AI	0.038***	0.064***	0.082***	0.110***	0.092**
	(3.17)	(2.95)	(2.84)	(3.13)	(2.16)
Automation	0.025*	0.035**	0.028	0.036	0.047
	(1.95)	(2.36)	(1.38)	(1.63)	(1.64)
Observations	146492	132119	120057	109327	99755
R-squared	0.075	0.101	0.107	0.114	0.122

Table 8
Continued

	Panel D: Cost Efficiency				
Automation Physical AI	-0.016	0.029	-0.185*	-0.266*	-0.276*
	(-0.15)	(0.21)	(-1.69)	(-1.75)	(-1.90)
Non-Automation Physical AI	0.012	-0.027	0.053	-0.165*	-0.008
	(0.15)	(-0.33)	(0.38)	(-1.76)	(-0.08)
Automation	-0.026	0.057	0.154	0.190	0.191
	(-0.22)	(0.94)	(1.38)	(1.39)	(1.52)
Observations	147888	133328	121099	110217	100535
R-squared	0.019	0.031	0.028	0.045	0.039
	Panel E: Profit				
Automation Physical AI	0.026**	0.024	0.044	0.047	0.068*
	(2.29)	(1.24)	(1.66)	(1.38)	(1.73)
Non-Automation Physical AI	0.044***	0.092***	0.119***	0.125***	0.136***
	(3.40)	(4.21)	(4.47)	(4.00)	(3.38)
Automation	0.027***	0.050***	0.057***	0.076***	0.078***
	(3.30)	(3.88)	(3.10)	(4.05)	(3.47)
Observations	151001	136828	124365	113339	103522
R-squared	0.091	0.110	0.117	0.123	0.133
Industry x Year FE	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year

Table 9: Product and Process Innovation as Physical AI Innovation

This table presents the relation between physical AI innovation, taking the form of product or process innovation, and firm-level outcomes. We present the coefficients β_τ for cumulative forward time periods $\tau \in \{1, \dots, 5\}$ from estimations of the following model:
$\ln(Y_{i,t+\tau}) - \ln(Y_{it}) = \alpha_0 + \beta_{1\tau} ProdPhysicalIAI_{it} + \beta_{2\tau} ProcessPhysicalAI_{it} + \beta_{3\tau} Process_{it} + \gamma_1 X'_{it} + \delta_{jt} + \epsilon_{it}$
where Y are firm outcomes represented in the Panel names; $ProdPhysicalAI$ is the fraction of the firm's patents that are AI-citing product patents; $ProcessPhysicalAI$ is the fraction of the firm's patents that are AI-citing process patents. *Process* is the fraction of a firm's patents that are process patents. Process patents are patents with more than 50% process claims. The rest are classified as product patents. The classification of claims is described in Appendix C. X is a vector of firm-time level control variables including log value of capital stock, the log number of employees, log idiosyncratic volatility, $\ln(1+Patents_{i,t-1})$, and $\ln(Y_{i,t-1})$; δ_{jt} are SIC2 times year fixed effects. Standard errors are double-clustered by firm and year. ***, **,* indicate significance at the 1%, 5%, and 10% respectively. All variable definitions are provided in Appendix A. The sample period is 1979 to 2024.

	t+1	t+2	t+3	t+4	t+5
			Panel A: Employment		
Product Physical AI	0.027***	0.042***	0.049**	0.059**	0.058*
	(2.91)	(2.83)	(2.63)	(2.42)	(1.90)
Process Physical AI	0.025**	0.054**	0.118***	0.111***	0.145***
	(2.04)	(2.26)	(3.75)	(2.97)	(3.34)
Process	0.011**	0.004	-0.003	-0.005	-0.016
	(2.34)	(0.56)	(-0.31)	(-0.40)	(-1.09)
Observations	164385	148238	134035	121597	110643
R-squared	0.057	0.073	0.083	0.092	0.101
			Panel B: Capital		
Product Physical AI	0.012	0.029	0.040*	0.034	0.021
	(0.99)	(1.58)	(1.75)	(1.09)	(0.54)
Process Physical AI	0.029	0.019	0.065	0.060	0.078
	(1.57)	(0.61)	(1.60)	(1.22)	(1.39)
Process	0.017**	0.031***	0.022	0.034*	0.027
	(2.14)	(3.12)	(1.47)	(2.01)	(1.27)
Observations	167201	150968	136597	123964	112783
R-squared	0.078	0.099	0.115	0.124	0.132
			Panel C: Output		
Product Physical AI	0.053***	0.056***	0.076***	0.095***	0.102***
	(5.33)	(3.24)	(3.23)	(3.56)	(2.91)
Process Physical AI	0.006	0.060*	0.095***	0.127***	0.136***
	(0.29)	(1.97)	(2.88)	(3.23)	(2.83)
Process	0.009	0.023*	0.017	0.006	0.018
	(1.01)	(1.97)	(0.98)	(0.29)	(0.75)
Observations	146492	132119	120057	109327	99755
R-squared	0.075	0.101	0.107	0.114	0.122

Table 9 (continued)

Panel D: Cost Efficiency					
Product Physical AI	-0.094*	0.054	-0.044	-0.130*	-0.039
	(-1.93)	(0.49)	(-0.44)	(-1.83)	(-0.58)
Process Physical AI	0.044	-0.084	0.090	-0.203	-0.168
	(0.31)	(-0.76)	(0.62)	(-1.32)	(-1.04)
Process	0.059	0.108	-0.007	0.124	0.136
	(0.64)	(1.12)	(-0.08)	(1.19)	(1.32)
Observations	147888	133328	121099	110217	100535
R-squared	0.019	0.031	0.028	0.045	0.039
Panel E: Profit					
Product Physical AI	0.056***	0.087***	0.091***	0.122***	0.133***
	(5.87)	(5.12)	(3.83)	(4.74)	(4.24)
Process Physical AI	0.023	0.056*	0.118***	0.107***	0.136***
	(1.35)	(1.97)	(3.48)	(2.73)	(3.03)
Process	0.020**	0.028**	0.023*	0.026	0.022
	(2.35)	(2.26)	(1.69)	(1.42)	(1.21)
Observations	151001	136828	124365	113339	103522
R-squared	0.091	0.110	0.117	0.123	0.133
Industry x Year FE	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year

APPENDIX A

Table A1. Variable Description

Variable	Description
ξ	Kogan et al (2017) value of a patent
σ	Idiosyncratic volatility in the patent grant year
Physical AI	The percentage of patents in a firm-year that cite an AI patent
AI Focal Patent	Indicator variable equal to 1 if focal patent is an AI patent identified as a patent with IPC classification G06N3, G06N5, G06N7, G06N10, G06F15, or G06K9. See Table A2 for a description of these classifications.
AI-Hiring	Babina et al. (2024) measure of AI employment
Automation	Fraction of a firm's patents that are automation patents. Appendix C.1. describes the classification of automation patents
Backward	Natural log of the number patents cited by the focal patent (i.e. backward citations)
Capital stock	COMPUTAT: ppent, deflated by the NIPA price of equipment
COGS/Sales	COMPUTAT: cogs/(sale+change in inventory)
Forward	Truncation bias adjusted forward citations calculated by dividing the number of forward citations a patent receives by the mean number of forward citations received by patents granted in the same year
Employment (Emp)	Number of employees
Exogenous AI	Number of AI patents granted to universities and federal labs located in the firm's commuting zone
$I_{PhysicalAI}$	Indicator variable equal to one for firm-years with at least one patent that cites an AI patent, and zero otherwise
I_{AIcite}	Indicator variable taking the value 1 if the patent cites at least one AI patent. AI patents are patents with IPC classification G06N3, G06N5, G06N7, G06N10, G06F15, or G06K9. See Table A2 for a description of these classifications.
$I_{Automation}$	Indicator variable equal to 1 if focal patent is an automation patent identified based on textual analysis of patent description. Appendix C.1. describes the classification of automation patents.
$I_{Process}$	Indicator variable that takes the value 1 if the more than 50% of the focal patent's claims are process claims. Appendix C.2. describes the classification of process claims.
Patents	Number of patents
Market cap	Shares outstanding times share price one day prior to patent grant date
Output	COMPUSTAT: sale plus change in inventories (invt) deflated by the CPI
PctAIcite	Percentage of cited patents that are AI patents
Process	Fraction of a firm's patents with more than 50% process claims. Appendix C.2. describes the classification of process claims.
Profit	Gross profit calculated as sales minus cost of goods sold
Revenue TFP	We follow Babina et. al (2024) and estimate Revenue TFP as the residual from regressing log real sales on log employment and log lagged capital, controlling for firm fixed effects and year fixed effects. The regression is estimated using OLS separately for each 2-digit SIC industry. The capital stock is constructed using the perpetual inventory method with a depreciation rate of 15% ($\delta = 0.15$). Real sales are deflated using the GDP deflator; real investment for the perpetual inventory method uses the non-residential investment price deflator

Table A2. Description of Artificial Intelligence (AI) IPC Classes

IPC Class	Description
G06N3	Computing arrangements based on biological models
G06N5	Computing arrangements using knowledge-based models
G06N7	Computing arrangements based on specific mathematical models
G06N10	Quantum computing, i.e. information processing based on quantum-mechanical phenomena.
G06F15	Digital computers in general; Data processing equipment
G06K9	Methods or arrangements for recognizing patterns

Table A3. Description of Information Technology (IT) IPC Classes

IPC Class	Description
G06F	Electric digital data processing
G06Q	Data processing systems or methods, specially adapted for administrative, commercial, financial, managerial, supervisory or forecasting purposes.
G06N	Computing arrangements based on specific computational models
G06K	Recognition of data; Presentation of data; Record carriers; Handling record carriers
G06T	Image data processing or generation, in general
H04L	Transmission of digital information, e.g. telegraphic communication
H04W	Wireless communication networks
H04M	Telephonic communications

APPENDIX B

Data From European Patent Office

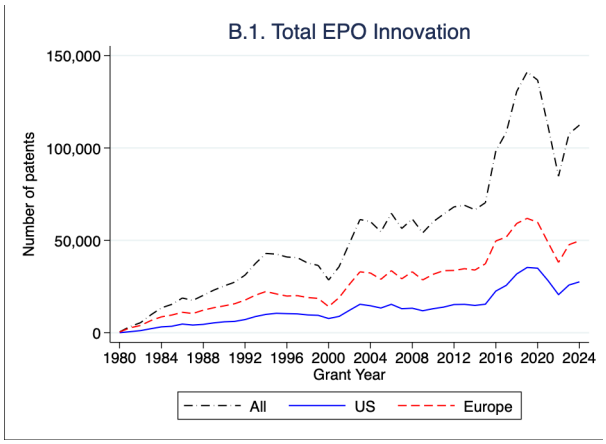
In Section IIa, we demonstrate that within the USPTO data, the U.S.–Europe gap in AI-absorbing innovation is widening, suggesting that U.S. firms are increasingly outperforming their European counterparts in integrating AI into the broader economy. To address concerns that U.S. dominance in AI-absorbing innovation merely reflects a home-market bias in USPTO data rather than a structural economic advantage, we examine patents issued by the European Patent Office (EPO). We extract just under 2.5 million patent grants spanning 1980–2024 from the EPO's Technology Intelligence Platform (TIP) and plot total patents and total non-IT patents by grant year in Figures B.1 and B.2, respectively. In these baselines, the well-documented home bias in patenting is highly evident—European firms maintain a substantial lead over U.S. firms in overall EPO grants (Criscuolo, 2006; Bacchiocchi and Montobbio, 2010).

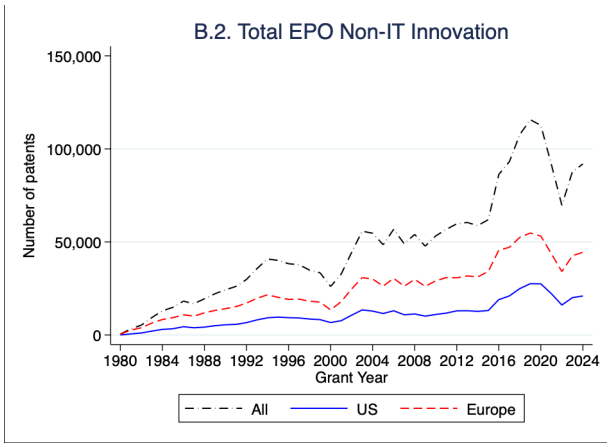
If U.S. dominance in AI and AI-absorbing patents within the USPTO data (Figure 2) were primarily an artifact of this home-market bias, one would expect European firms to exhibit a parallel dominance within the EPO data. We find this is not the case. Figures B.3 and B.4 depict the counts of AI patents and AI-absorbing non-IT patents issued by the EPO. Figure B.3 shows that U.S. firms marginally lead European firms in absolute AI patent grants, while Figure B.4 demonstrates that U.S. firms are broadly comparable to European firms in AI-absorbing grants. The fact that European firms fail to command a lead in AI-absorbing innovation even within their home patent office strongly reinforces our primary finding: U.S. strength in AI assimilation is an economically meaningful phenomenon, not an artifact of geographic data selection.

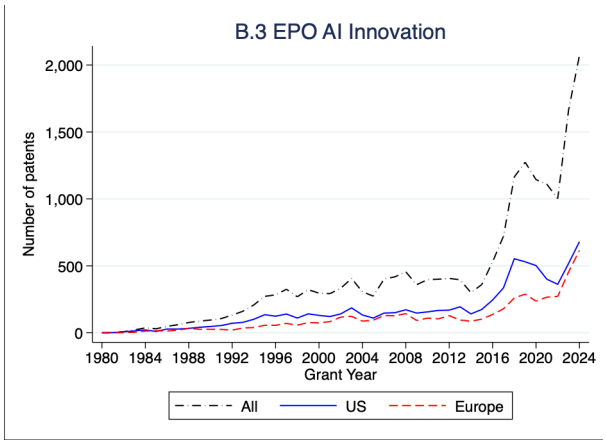
Figure B: EPO Patents

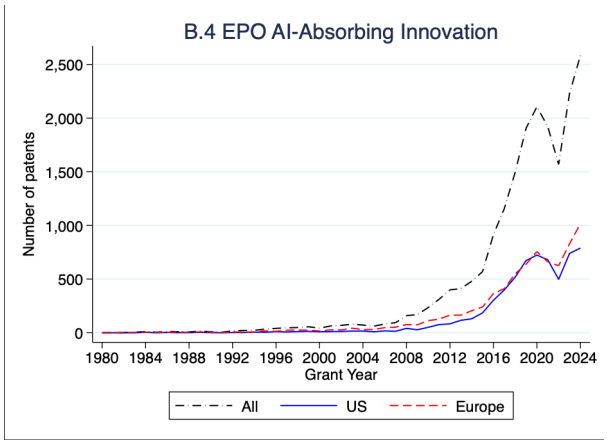
Figure B.1 depicts the total number EPO patents over the period 1980 to 2024 and the number of EPO patents granted to U.S firms and European firms respectively. Figure B.2. depicts the total number non-IT patents granted by EPO over the period 1980 to 2024 and the number of non-IT EPO patents granted to U.S firms and European firms respectively. Non-IT patents are patents *not* in the following information technology IPC classes: G06F, G06Q, G06N, G06K, G06T, H04L, H04W, H04M.

Figure B.3. depicts the number of Artificial Intelligence (AI) patents granted by EPO, and the number of AI patents granted to U.S. and European firms respectively. Patents with IPC class G06N3, G06N5, G06N7, G06N10, G06F15, and G06K9 are classified as AI patents. Figure B.4 depicts AI-absorbing EPO patents, defined as non-IT patents that cite AI patents.









Appendix C: Classification of Automation Patents and Process Patents

C.1. Classification of automation patents

We machine-read patent descriptions and classify a patent as an 'automation' patent if it meets any one of the following four criteria:

1. One or more appearances of the following words:
   - automatic or automatically
   - automate, automation, or automatization
   - mechanize or mechanization
   - robot, robotic, robotize, or robotization
   - labor-saving
2. To identify computer-assisted manufacturing: One or more appearances of the following bulleted list of words **and** the word *machine*, *manufacturing*, *equipment*, or *apparatus*
   - computer aided or computer-aided
   - computer assisted or computer-assisted
   - computer supported or computer-supported
3. To identify computer numerical control (cnc) patents which indicate automated control of machine tools: One or more appearances of the following words
   - numerically controlled or numerically-controlled
   - numeric control or numeric-control
   - cnc
4. To identify self-operated equipment: One or more appearances of the word 'self-' **and** the word *machine*, *manufacturing*, *equipment*, or *apparatus*

C.2. Classification of process claims

We follow Bena and Simintzi (2022, 2025) and take advantage of the standardized vocabulary with stilted legalistic terms used for process patents. Process patent claims often contain words such as 'method of', 'process for', to describe the steps or procedures involved in the invention. We classify a claim as process-oriented when its opening phrase matches the following pattern: an optional article ("a," "an," or "the"), followed by up to five optional descriptive modifiers, and then the keywords *method* or *process*. The expression also allows for hyphenated modifiers and flexible punctuation or spacing immediately following the keyword. Formally, this flexibly captures claim openings such as

- "a method…"
- "an improved manufacturing process…"
- "the computer-implemented method…," etc.

The matching pattern excludes occurrences in which *method* or *process* appears later in the sentence rather than as the claim's principal grammatical subject. Anchoring the pattern to the start of the claim text ensures that the classifier targets the conventional legal structure of process claims rather than incidental keyword usage elsewhere in the document.