Physical AI Innovation*

Adair Morse
University of California - Berkeley
& NBER

Amrita Nain
University of Iowa

November 2026

Abstract

We study physical AI innovation: innovation outside the information technology (IT) sector that incorporates artificial intelligence (AI). Physical AI patents are nearly as prevalent as AI patents, command excess financial market value, and reflect a persistent transatlantic gap. Tracing effects through production factors and outcomes and instrumenting with laboratory AI knowledge (universities and federal), we show that physical AI innovation increases annual growth in employment and capital stock by 1.5% and 3.1%, respectively. Under a simple Cobb-Douglas parameterization, the growth in capital and labor implies output growth of 2.0%, close to the 2.1% we estimate. The benefits of physical AI innovation also flow downstream: customers experience growth in employment, output, and TFP, albeit with a lag. We find no average effect of physical AI innovation on cost efficiency. One important exception arises: physical AI patents that are also automation patents – which we term *smart automation* – are associated with lower costs per unit of output, an effect concentrated in industries with high abstract/cognitive task intensity.

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

* Author contacts: adair@berkeley.edu and amrita-nain@uiowa.edu. We thank Renping Li, Gerald Hoberg, and seminar participants at the University of Iowa, University of Richmond, University of Mannheim. We also thank participants of the Conference on Technology, Financial Services, and Evolving Societal Needs at the University of Gothenburg, and the CREDIT 2026 Conference on Artificial Intelligence, Financial Innovation and Credi Risk: Sustainability, Stability and Responsibility in a Data-Driven Financial System. This paper was previously titled AI-Absorbing Innovation.

I. Introduction

Artificial intelligence (AI) is a general-purpose technology (Brynjolfsson, Rock, and Syverson, 2021), and much like the steam engine, its value comes from downstream use in production across the real economy. As noted in the foundational work of Bresnahan and Trajtenberg (1995), general-purpose technologies are characterized by "*the potential for pervasive use in a wide range of sectors*" and can lead to "*innovation complementarities*" in other sectors. In Silicon Valley parlance, the deployment of AI into non-technology sectors is increasingly referred to as physical AI.

A growing literature studies the effect of AI on firm outcomes using AI input measures such hiring of AI experts, task-level exposure to AI, investments in data centers, and the use of generative AI agents.[1] However, we know little about how AI technologies are absorbed into innovation across the economy and the downstream implications of this absorption. One reason for this gap is that the diffusion of AI knowledge and the downstream impact of AI are difficult to observe and measure (Korinek and McKelvey, 2026).

We fill this gap by studying the effect of complementary non-IT innovation that used AI, i.e. physical AI innovation, on the production outcomes of firms and their customers. Conceptually, physical AI is AI that operates on or through the physical world: collecting information from physical environments, optimizing physical processes, and adapting to changing conditions in real time. We operationalize this concept by developing a patent-based measure of physical AI innovation: non-IT patents that cite prior AI patents. Our backward citation-based measure builds on 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). This measure lets us study whether the value of AI emerges when combined with complementary innovation in other sectors, as would be expected for a general-purpose technology (Bresnahan and Trajtenberg, 1995; Brynjolfsson, Rock, and Syverson, 2021).[2]

[1] AI hiring, job postings, or task exposure are used in Acemoglu, Autor, Hazell, and Restrepo (2022); Babina, Fedyk, He, and Hodson (2024); Hampole, Papanikolaou, Schmidt, and Seegmiller (2025); Brynjolfsson, Chandar, and Chen (2025); and Albanesi, Dias da Silva, Jimeno, Lamo, and Wabitsch (2025). Eisfeldt, Schubert, and Zhang (2023) combine workforce exposure with the release of ChatGPT. Acharya, Shao, and Xu (2026) use installations of AI products, and Brynjolfsson, Li, and Raymond (2025) study the rollout of a generative AI assistant at a firm. Alvarez, Argente, Chow, and Van Patten (2026) and Yue and Zeng (2026) use data center investment.

[2] See also Teece (1986); Bresnahan, Brynjolfsson, and Hitt (2002); and Fedyk, Gomes, Mihet, and Rishabh (2026) on how complementary assets and capabilities affect the value firms obtain from technology.

The absorption of AI technologies into the non-IT sectors is important to study because 80% of USPTO patents lie outside IT patent classes, and physical AI patents are of the same order of magnitude as AI patents themselves. This issue also has macro relevance. The growth differential between the United States and Europe has been attributed in part to the U.S. advantage in producing frontier IT innovation.[3] However, that attribution overlooks the U.S. advantage in physical AI innovation.

We show that physical AI innovation is both prevalent and economically valuable, and that the U.S. share of physical AI patents exceeds Europe's by about 50 percentage points. We study the effect of physical AI innovation on the factors of production (labor, capital), total factor productivity (TFP), output, and cost efficiency for focal firms (firms outside the software and communications sector) and their corporate customers. Ex ante, the effects are ambiguous: AI could substitute for labor, raise productivity, or expand what firms produce. We document positive effects of physical AI for focal firms and, at longer horizons, for their customers. The gains operate through the expansion of inputs and output, not through cost reductions. To our knowledge, we provide the first evidence on how AI knowledge diffuses across sectors and on the benefits that reach downstream firms.

Our work builds on two strands of research: the firm-level literature examining AI's economic effects through labor substitution, task augmentation, and productivity (see Footnote 1) and the macroeconomic literature viewing AI as a technology shock to the innovation process itself, with consequences for aggregate growth and market value.[4] We fit most closely with Babina et al. (2024), Fedyk, Gomes, Mihet, Rishabh (2026), and Argente, Baslandze, Hanley, and Moreira (2025), each of which bridges the two literatures by examining how new AI and data technologies affect firm-level innovation and growth. Our angle is distinct because we focus not on AI innovation itself, but on non-IT inventions that incorporate AI advances into the physical economy. In doing so, we contribute to the growing body of research studying AI's consequences for firm growth, productivity, and asset dynamics.[5]

---

[3] 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)

[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.

[5] E.g., Kelly, 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 al., 2026; and Baslandze et al. (2026).

To fix ideas, we provide an example of a physical AI invention: a 2021 microsurgery device patent (US 10,895,742).[6] 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 AI advances in biometrics and facial-recognition technologies 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 a few points. First, AI's influence on invention reaches well beyond the technology patents that originate it. 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, in the spirit of the innovational complementarities of Bresnahan and Trajtenberg (1995). Third, the benefits can pass downstream: a hospital that uses the device may deliver higher-quality care and train its surgeons more effectively. Physical AI is also broader than the common conception that it is synonymous with robots and autonomous control. We view physical AI as AI that operates on or through the physical world: collecting information from physical environments, optimizing physical processes, and adapting dynamically to changing real-world conditions. These capabilities appear in medical devices, as illustrated here, but also in other production instruments and devices, logistics and financial services, and systems in agriculture, energy, and infrastructure.

Before turning to production implications, we characterize physical AI. Our analysis uses all patents granted by the United States Patent and Trademark Office (USPTO) to publicly traded firms between 1979 and 2024. Among non-technology patents issued to U.S. firms, 4% are physical AI. Further, the physical AI patents have excess market valuation. In the subset of patenting firms in the non-IT sector, 7% have at least one physical AI patent. Using the patent value measure of Kogan, Papanikolaou, Seru, and Stoffman (2017) that looks within the same firm-year and technology class, non-IT patents that cite AI patents command a 40-basis-point valuation premium relative to other non-IT patents. Looking within the subset of IT patents, we find that AI-citing patents have no such premium. This contrast is consistent with AI being a

[6] 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.

general-purpose technology, whose value emerges when it is combined with domain-specific knowledge outside the technology space.

Before delving into our analysis, we ask whether physical AI innovation simply relabels existing measures. We show that process patents and automation patents explain little of the variation in physical AI, with partial R-squared values around 0.11. At the firm level, the overlap with the AI-skilled hiring measure of Babina et al. (2024) is also low. This is not surprising, because the two capture different things: hiring is an input into AI efforts, while citations to AI patents are evidence that AI knowledge has been incorporated into non-AI innovation. AI hiring may not capture the full range of real-economy innovation that incorporates AI only partially.[7] Conversely, citation-based absorption may miss firms' investing in AI expertise that focuses on implementation instead of patenting. The two measures are likely complements, not substitutes.

Firms that engage in physical AI innovation may have more profitable innovation opportunities and a greater stock of knowledge capital. Therefore, to estimate the effect of physical AI innovation on the outcomes of focal firms and their customers, we use an instrumental variables implementation of local projections (LP-IV), following Jordà (2005) and Ramey (2016). We exploit localized knowledge spillovers, instrumenting firms' physical AI innovation with the local stock of AI patents produced by universities and federal laboratories in the same commuting zone as the firm's inventors, building on Jaffe, Trajtenberg, and Henderson (1993) and Helper, Makan, and Shoag (2025). We refer to our instrument as Laboratory AI. The exclusion restriction requires that, 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. Laboratory AI must affect outcomes only through the absorption of AI knowledge.

We address three threats to this exclusion restriction. First, Laboratory AI could proxy for local economic conditions. Because about 40% of patents have inventors outside the firm's headquarters commuting zone, we can include headquarters commuting zone times year fixed effects and compare firms headquartered in the same place with inventors located in different commuting zones. Because inventors may instead be stationed at a firm's primary operating locations, we also replace these with primary operating state times year fixed effects, and the

[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.

results are robust. Second, firms with stronger growth prospects may move inventors to AI hubs in anticipation of expansion. Firm fixed effects absorb static location choices, and we address dynamic location choice by dropping the five years after a firm first places inventors in a commuting zone. Third, universities and federal laboratories also produce non-AI research, so Laboratory AI could capture general research spillovers. A placebo test finds no relationship between Laboratory AI and the outcomes of innovative firms that do not engage in physical AI innovation.

Whether physical AI innovation should raise employment is not obvious. If it substitutes for labor, as concerns about AI and automation suggest (Acemoglu et al., 2022; Hampole et al., 2025), firms could grow output while shedding workers or capital intensity could rise without employment. If it instead expands what firms can invent and produce, labor and capital should grow together. We find the latter. For focal firms, we find that a one-standard-deviation increase in physical AI innovation raises the annualized growth rate over a two-year horizon by about 1.5% in employment, 3.1% in capital, and 2.1% in output. The employment and output effects are concentrated in the first two years. The output growth is consistent with the growth in inputs alone: under a Cobb-Douglas production function, the growth in capital and labor implies output growth of 2.0%, close to the 2.1% we estimate. While we observe some evidence of an increase in TFP for focal firms engaged in physical AI innovation, the findings are not robust in more stringent specifications. Moreover, we find no evidence of an average reduction in costs associated with physical AI innovation.

For customers, we observe positive effects with a lag. We find no effect in the first two years following upstream physical AI innovation by the focal firm. Employment, output, and TFP then rise from the third year, with cumulative effects peaking around year four. A one-standard-deviation increase in the focal firm's physical AI innovation raises customers' annualized growth rate over four years by about 0.4% in employment, 0.8% in output, and 0.4% in TFP. The lag is plausible: customers may be locked into existing contracts, processes, or products, and must observe upstream changes and adapt their own operations before the benefits appear. Unlike for focal firms, we find a productivity effect: customers appear to use their existing labor and capital more efficiently.

Overall, our evidence suggests that firms invest in physical AI innovation because of expansion opportunities, and not to achieve cost reductions. These results are consistent with,

albeit through a different mechanism than, the innovation outcome effects documented in Babina et al. (2024). One important exception emerges when we decompose physical AI into automation and non-automation components. Approximately 35% of physical AI patents also qualify as automation patents; we term this type of innovation as "S*mart Automation"*. While non-automation physical AI shows no cost efficiency effect, Smart Automation has a strong and distinct cost-reducing effect: at a five-year horizon, a one-standard-deviation increase in Smart Automation predicts an annualized reduction in COGS/output of 0.0064.

As a final step, we ask whether the cost-reducing effect of Smart Automation is concentrated in industries where AI capabilities can replace human adaptive reasoning and cognitive skills. Following Autor, Levy, and Murnane (2003) and Autor, Katz, and Kearney (2006), we classify industries by their abstract/cognitive task intensity at the 3-digit SIC level. Consistent with intuition and the inferences in Hampole et al. (2025), the cost efficiency effect of Smart Automation is entirely driven by firms in industries with high concentration of tasks requiring abstract and cognitive thinking.

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 incorporation into innovation outside the IT sector. Existing work studies either the labor market effects of AI adoption or the role of AI innovation in driving technological progress and growth (e.g., Acemoglu et al., 2022; Babina et al., 2024; Hampole et al., 2025). We introduce physical AI innovation as a distinct and empirically important category, document its scale, its market valuation premium, and its pronounced U.S.-Europe gap, and provide the first systematic evidence on its consequences for firm-level production for both the inventing firms and their customers.

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 automation-aimed innovation activities of non-IT firms. We

label such innovation as "smart automation". 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. We study smart automation from the innovation side, finding that it predicts cost efficiency gains – a channel distinct from the revenue-expansion effect that characterizes non-automation physical AI.

More broadly, our paper contributes to the literature on absorptive capacity and knowledge spillovers (Cohen and Levinthal, 1990; 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.

Finally, we present some evidence building on the process and production innovation literature. 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.[8]

The rest of the paper is organized as follows. Section II discusses the patent data and characterizes physical AI innovation. Section III presents our methodology, and Section IV presents the main results relating to real-economy inputs and outputs. Section V concludes.

[8] 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.

II. Data and Statistics

II.a Patents & Physical AI Innovation: Measurement

We download 8.4 million utility patents granted by the U.S. Patent and Trademark Office (USPTO) between 1979 and 2024 from USPTO PatentsView. 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.[9] 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 complementary assets in these sectors absorb and translate AI knowledge into value and employment creation is of first-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 because the much-discussed innovation gap between them may contribute to their diverging economic performance. For example, 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). In Figures 1.C and 1.D, we plot the innovation gap —the U.S. share of patents granted minus the European share — for all patents and non-IT patents respectively. The U.S. advantage peaked in the mid-1990s and has since narrowed in both categories, but the contraction is most pronounced in non-IT sectors. The second takeaway from the graphs is that Europe has been closing the gap in non-IT innovation.

In Figures 2.A and 2.B, we plot total AI patents granted and total AI-absorbing patents, respectively.[10] Following the citation-based convention for measuring knowledge absorption, we define AI-absorbing patents (interchangeably, physical AI innovation) as non-IT patents that cite a prior AI patent. These are inventions in sectors such as manufacturing, aerospace, energy, and

[9] See appendix Table A2 for the classification of IT patents.
[10] See appendix Table A3 for the classification of AI patents.

agriculture that incorporate AI technologies.[11] Because the commercial value of such inventions accrues in the application domain, the patents are classified by that domain: an aerospace firm's AI-enabled guidance system is coded as an aerospace patent, not an AI patent.

Figures 2.A and 2.B provide the third pattern we wish to emphasize: patenting that absorbs AI into the physical economy is of the same order of magnitude 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 present our fourth graphical finding: the U.S.–Europe gaps in AI innovation and physical AI innovation. The magnitude of the transatlantic AI inventiveness gaps are both stable but very large (approximately at 0.5) relative to Figures 1.C and 1.D. The U.S. advantage over Europe is not simply in producing AI, but in absorbing AI into the broader economy.[12] To the extent that AI-absorbing innovation drives firm production and market valuation, this pattern carries consequences for both economies' trajectories and employment.

II.b Characterizing Physical AI Innovation: Industry Mix

To shed light on which sectors are absorbing AI into the physical economy, we focus on patents granted to publicly traded U.S. firms. We obtain patent data with matched CRSP permanent numbers from the Kogan et al. (2017) data library, restrict the sample to 1979–2024, and match to Compustat. We then drop firms in software and communications (SIC codes 357, 481, 482, 484, 489, 737), and for the remaining physical-economy firms, count the number of firms in each 2-digit SIC code with at least one non-IT patent citing an AI patent.

Table 1 reports the top 2-digit SIC codes for firms engaged in physical AI innovation. The breadth is striking. The most prominent cluster involves instruments and equipment — a category

[11] In some forums, physical AI refers to robotics and automation only, but this is not the common usage definition emerging in Silicon Valley.

[12] 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. Furthermore, Appendix B presents the graphs with European Patent Office (EPO) grants to check that the U.S. advantage does not simply reflect home-market bias in USPTO data.

spanning healthcare, transportation, manufacturing, construction, and energy. A second cluster comprises service industries: financial services and insurance, energy services, printing, and engineering. A third encompasses product-oriented firms in chemicals, natural resources, and industrial goods. Physical AI innovation is not concentrated in a narrow slice of the economy; it is pervasive.

II.c Characterizing Physical AI Innovation: Relation to Prior Literature Measures

Before proceeding, we ask whether our measure identifies a genuinely new concept or simply relabels existing ones. Our physical AI measure might be highly correlated with automation,[13] process patenting,[14] or AI hiring.[15] We classify patents into process and product patents following Bena et al. (2022) and Bena and Simintzi (2025) using textual analysis (see Appendix C2 for a description of process patent classification). We identify automation patents using natural language processing applied to each patent's textual description, detailed in Appendix C1. Finally, we obtain firm-level AI hiring data from Babina et al. (2024).

As shown in appendix Table A4, panel A, among non-IT patents, 4.1% are AI citing, 13.9% are automation, and 28.4% are process patents. Automation and process patenting are considerably broader concepts than Physical AI innovation. Panels B and C report the partial R-squared from regressing physical AI innovation on automation, process patenting, and AI hiring at the patent level (Panel B) and firm level (Panel C), with and without SIC4 × year fixed effects. The explanatory power is modest throughout: the highest partial R-squared is 12%, from automation's correlation with firm-level physical AI variation.

Our measure also differs from the existing AI-employment nexus literature in three ways

[13] Autor et al. (2003); Acemoglu and Restrepo (2019) document the nexus between automation and labor-saving efficiency. The linkage of AI effects goes beyond just employment. 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 (Autor and Dorn, 2013; Berman, Bound, and Griliches, 1994; Berman, Bound, and Machin, 1998; Acemoglu, 1998; Autor, Katz, and Krueger, 1998).

[14] 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. Closely related to our work, 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.

[15] Babina et al. (2024), who find that firms' investments in AI hiring stimulate product innovation while having little effect on process innovation.

that are central to our contribution: we focus on innovation activity rather than AI adoption measured by 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.[16]

II.d Characterizing Physical AI Innovation: Patent-Level Statistics

We create a patent-level I_{AIcite}, an indicator variable that takes the value of one if a patent makes at least one backward citation to an AI patent. Panel A of Table 2 presents the summary statistics for 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 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 16% of the patents cite at least one 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 0.041; about 4% of non-IT patents cite at least one AI patent. We note that the rest of the patent characteristics such as forward and backward citations, firm size, volatility are similar across Panel A and B.

To further characterize physical AI innovation, we compare AI-citing patents to other patents in terms of economic value using the patent market value -- the (ξ) -- of Kogan et al. (2017). We estimate the descriptive correlation between AI absorption and value of a mean patent *p* granted to firm *i* as follows:

$$ln(\xi_p) = \alpha + \beta\, I_{AIcite_p} + \gamma_1 \ln Market\ Cap_p + \gamma_2 \ln Backward_p + \mu_{i,t} + \rho_k + \epsilon_p. \quad (1)$$

$ln\, Backward$ is the log number of patents cited by the patent, included because assimilation of any existing knowledge can be valuable. *Market Cap* is the firm's log market capitalization measured the day prior to patent grant date (patenting may increase with size). $\mu_{i,t}$ are firm-by-grant-year fixed effects which control for characteristics such as idiosyncratic volatility that are

[16] 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.

correlated with ξ (e.g., high-growth firms having more volatile returns and higher patenting value). Finally, we include technology class fixed effects, ρ_k, to force the comparison of AI assimilation to within technology comparables. As shown in the descriptive Table 3,

In the descriptive Table 3, Column 1 shows that, in a subsample of IT patents, the coefficient on I_{AIcite} is insignificant. That is, AI absorbing IT patents do not have higher market valuations than other IT patents. In contrast, within the sample of non-IT patents (Column 2), AI absorbing patents (i.e., physical AI citations) command higher economic value than other non-IT patents, significant at the 5% level. 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 relative to a patent that does not cite any AI.[17] Our finding suggests that a meaningful portion of AI's value emerges when it is combined with domain-specific knowledge from other sectors, reminiscent of Bresnahan and Trajtenberg (1995) and Zahra and George (2002)'s seminal work on technology value pathways and transformative capability in the synthesizing of knowledge across unconnected technologies.

II.e Characterizing Physical AI Innovation: Firm-level Statistics

At the firm level, we define *PhysicalAI* as the percentage of non-IT patents filed in a firm-year by non-IT firms – equivalently, the firm-year average of the patent-level indicator variable I_{AIcite}. This variable takes the value zero for firms with no AI-citing patents or no patents. We define non-IT firms as firms that are not in the software and communications sector.[18] For comparison, we also study a variable $I_{PhyscialAI}$ equal to one for firm-years with at least one patent that cites an AI patent, and zero otherwise.[19]

Table 4 summarizes *PhysicalAI* and $I_{PhysicalAI}$, as well as all of the production input and output variables and control variables. In the full sample, 1% of the stock of patents in a firm-year cite at least one AI patent, with an average of $I_{PhysicalAI}$ of 0.029, implying that about 3% of firm-years have at least one focal patent that cites an AI patent. These overall percentages are low partly because the majority of firms in the sample have no patents and, therefore, have zero values of

[17] Since the dependent variable is the natural log of ξ, and I_{AIcite} is an indicator variable, the coefficient 0.0039 in column 2 represents ($e^{0.0039}$-1)*100= 0.4% higher market value.
[18] Thus, PhysicalAI excludes IT patent classes as well as IT sectors. The excluded software and communications sector is defined as firms with the following three-digit SIC codes - 357, 481, 482, 484, 489, and 737.
[19] Our results are similar if we conducted analysis with $I_{PhysicalAI}$ rather than *PhysicalAI*. Our results are also robust if we exclude firms that have no patents in our sample period.

PhysicalAI and $I_{PhysicalAI}$. Within the subsample of non-IT firms that have at least one patent, we find that 7% engage in physical AI innovation over the sample period. Size also plays a role. If we focus on the top quartile of non-IT firms by size – those with market capitalization over $491 million, – 8.4% of firms have physical AI innovation across our time period. Given that the top size quartile in our sample accounts for $20 trillion in total market capitalization on average over the period and that physical AI is increasing over time, the implementation of physical AI is economically quite meaningful.

We are interested in studying the effect of Physical AI on the production outcomes of the focal firms as well as their customers. Publicly traded firms are required to disclose customers that account for more than 10% of revenue.[20] We extract customer names from Compustat segment files and keep all records that name a corporate customer over our sample period. We drop unnamed entities and government owned entities. The remaining names are matched to Compustat gvkey using a combination of exact name matching, fuzzy matching, and manual matching that accounts for name changes over time. We are able to match 30% of distinct corporate customer names, which account for 54% of named corporate customer disclosures. Most of the remaining customers are private, foreign or subsidiary firms that have no Compustat record. We further drop customers that have no CRSP record and link the remaining disclosed customers to our focal firm sample. We are able to match 2,667 unique customer firms to just over 5,000 unique focal firms. Panel B of Table 4 summarizes outcome variables and control variables for the customer firms over our sample period. The customer sample is considerably smaller than the focal firm sample because it is restricted to large customers: output, employment, and capital stock are noticeably larger for customer firms than for focal firms. In Panel B, we also summarize focal firm (i.e. supplier firm) Physical AI and Laboratory AI. The mean values of focal firm Physical AI and Laboratory AI in Panel B are comparable to the full sample means in Panel A.

III. Methodology: Laboratory AI IV

We examine both the input and output sides of production. For inputs, capital and labor are the natural production factors; we measure them using employment and fixed asset capital stock

[20] Under SFAS 131, firms must disclose revenue from any customer that accounts for 10% or more of total revenue. Although the standard does not require firms to identify the customers, many do, in part because SEC Regulation S-K requires firms to name customers whose loss would have a material adverse effect

from Compustat. For factor productivity, we apply Babina et al.'s (2024) revenue-based TFP specification. Specifically, we assume a Cobb-Douglas technology for 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:

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

Following Babina et al. (2024), we allow the factor intensity parameters to vary by sector s, and include firm and time t fixed effects. On the production outcome side, we look for effects in output (revenues plus the change in inventories) and cost efficiency (COGS divided by output).

Following Kogan et al. (2017), Kelly, Papanikolaou, Seru, and Taddy (2021) and Baslandze et al. (2025), a Local Projection (LP) methodology (Jordà, 2005) could estimate firm i's outcome response to *PhysicalAI* at horizon τ according to:

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

The dependent variable in this specification is the τ-year log change in Y, where Y is alternatively the number of employees, the capital stock, revenue TFP, output, and cost efficiency, with $\tau = 1, \dots, 5$ years. Controls in vector X_{it} include the lagged dependent variable, $Y_{i,t-1}$. firm size, risk dynamics, and the log of one plus the number of patents filed in the year as a measure of innovative capacity. Finally, the specification includes firm fixed effects μ_i and SIC2 times year fixed effects δ_{jt} to isolate the dynamic control group within industry-year, with standard errors double clustered by firm and year.

The problem with equation (2) is that firms engaging in physical AI innovation may already be on a different growth or productivity trajectory, in which case β_τ would not identify the causal effect of physical AI. Hence, following the LP-instrumental variable (LP-IV) methodology of Ramey (2016), we instead instrument a firm's *PhysicalAI* with the local stock of AI patents and AI-citing patents generated by universities and federal laboratories in the same commuting zone (CZ)**.** This strategy builds off a large literature on absorptive capacity spillovers and scientific knowledge.[21] Universities and federal laboratories generate frontier scientific knowledge that is plausibly exogenous to individual firms' production decisions, particularly among the non-IT firms

[21] E.g., Jaffe, Trajtenberg, Henderson (1993); Azoulay, Graff Zivin, and Wang (2010); Balsmeier, Fleming, and Lück (2023); Helper et al. (2025); and Arts, Fleming, and Veretennik (2026)

in our sample. A higher local AI knowledge stock (*Laboratory AI*) from universities and laboratories increases a firm's exposure to AI knowledge through localized spillovers.

The exclusion restriction requires that, conditional on controls, *Laboratory AI* is unrelated to firm outcomes except through the firm's AI knowledge absorption. The exclusion restriction would be violated if Laboratory AI raised firm outcomes through the local economy, for example through demand, agglomeration, or local market conditions where the firm is located. The construction of our instrument enables us to control for local market conditions.

We calculate the local AI knowledge stock for patent p filed in year t as the mean annual sum of AI patents plus AI-citing patents granted to universities and federal laboratories located in the same commuting zone (CZ) as the inventor of patent p, with the mean calculated over the five years from *t-5* to *t-1*.[22] For just over 40% of patents, the inventor is located in a different CZ than the assignee firm's headquarters. This variation allows us to include headquarter CZ times year fixed effects, $\psi_{CZ,t}$, which absorb local conditions common to firms headquarters in the same CZ in a given year. Identification then comes from differences across firms headquarters in the same CZ. To illustrate, consider two farm machinery firms, A and B, headquartered in Kansas City with the same level of patenting. Firm A has inventors in the San Francisco Bay Area, where Laboratory AI is high, and Firm B does not. Firm A, therefore, has higher Laboratory AI, even though both firms face the same local economy at headquarters. The size of the difference depends on the share of each firm's innovation invented in the Bay Area.[23]

The exclusion restriction also requires that firms' choice of inventor locations is unrelated to their outcomes. Consider why Firm A has inventors in the Bay Area while Firm B does not. Firm A may have opened an R&D office there to gain access to the knowledge spillovers of an AI hub. If this choice is static, in that Firm A has always had inventors in the Bay Area, firm fixed effects, μ_i, absorb it. Of greater concern is dynamic choice. Firms with stronger growth prospects may move inventors to high *Laboratory AI* locations in anticipation of expansion. Firm-fixed effects cannot absorb this, and *Laboratory AI* would then proxy for time-varying unobserved firm

[22] If a patent has multiple inventors, we use the highest value of local AI stock among its inventors.

[23] Inventors may be stationed at a firm's primary operating locations instead of its headquarters, in which case Laboratory AI could reflect local conditions that headquarters CZ × year fixed effects do not absorb. In untabulated analysis, we re-estimate the specification using primary operating state × year fixed effects instead. Following the text-analysis approach of Garcia and Norli (2012), we define a firm's primary operating state as the state most frequently referenced in Items 1, 2, 6, and 7 of its 10-K.

characteristics rather than exposure to AI knowledge.[24] To address dynamic selection, we run "incumbent" specifications that drop the year a firm first places an inventor in a commuting zone, along with any of the following four years in which the firm has an inventor there, since these are the years in which a decision to relocate inventors is most likely to reflect anticipated growth.

Another threat to the exclusion restriction is that universities and federal laboratories also produce substantial non-AI research, so Laboratory AI could capture general research spillovers unrelated to AI knowledge. If so, Laboratory AI would predict outcomes in the subset of innovative firms that do not engage in physical AI innovation. We show that it does not (Appendix Table A5).

Figure 3 maps Laboratory AI across U.S. commuting zones. We highlight three commuting zones, circled in red, to show the different sources of this variation: universities, federal laboratories, or both. The San Francisco Bay Area is among the most AI-intensive commuting zones in the country. Its AI knowledge comes from universities such as UC Berkeley and UCSF and from two major Department of Energy (DOE) laboratories, Lawrence Livermore and Lawrence Berkeley. Iowa City / Davenport has moderate Laboratory AI, driven almost entirely by the University of Iowa, although the U.S. Geological Survey's Central Midwest Water Science Center is also located there. The Washington, D.C. metro area (the "DMV," for the District, Maryland, and Virginia) is different again. It hosts notable universities, including Georgetown, the University of Maryland, George Washington, and George Mason, but its Laboratory AI draws more heavily on federal institutions: NIH, NIST, NASA, and a dense concentration of DOE and Department of Defense facilities.

The distribution of Laboratory AI necessitates two further checks. First, Laboratory AI is highly concentrated in a few places: the San Francisco Bay Area (excluding the South Bay), the South Bay (San Jose and Silicon Valley), Austin-Round Rock, greater Boston, and the DMV. Because these five hubs account for a large share of Laboratory AI, we drop firm-years in which the firm filed at least one patent by an inventor located in one of them, to show that our results do not depend on a handful of locations (Appendix Table A6). Second, firms without patents have no inventor location and therefore a Laboratory AI value of zero. The results could then reflect a comparison of non-inventing firms with inventing firms, not the effect of AI knowledge. Figure 3

[24] 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. Yet, the correlation between firm-level AI hiring (from Babina et al (2024) and *PhysicalAI* is very low.

shows that no commuting zone has zero Laboratory AI, so the zeros arise from firms without inventors and not from places without AI research. We, therefore, also estimate the specification within innovative firms only, defined as firms with at least one patent over the sample period (Appendix Table A7).

IV. Results

Following the split introduced in the data section, we present results separately for focal firms and their customers. We think of AI patents as the 'steam engine,' focal firms as the innovators that build it into their own products and processes, and customers as the downstream users of those inventions.

Before delving into the second stage results, we pause to consider first-stage estimations for relevance. Table 5 reports first stage estimates for the three main subsequent analyses. Column (1) considers the main IV estimate for the focal firms. Column (2) also consider the focal firms, but the sample is restricted to be the incumbent analysis, which excludes the first year n in which a firm has an inventor in a given commuting zone, as well as any year from $n+1$ to $n+4$ in which the firm has an inventor in that same commuting zone. Finally, Column (3) reports the first stage for the customer sample, in which the dependent variable is the upstream focal firm's physical AI innovation and the instrument is the Laboratory AI of the upstream firm. Each row block is a separate estimation for a forward horizon $\tau = 1, \ldots, 5$ years. The second-stage outcome that defines each first-stage sample, not shown in Table 5, is the change in employment over τ years relative to year t.

In Columns (1) and (2), we find that *Laboratory AI* robustly positively predicts *Physical AI,* with Kleibergen-Paap F-statistics higher than 117 across all horizons, well above conventional weak-instrument thresholds. In Column (3), we instrument *UpstreamPhysicalAI* with *Laboratory AI*. Because we are limiting the sample to observable customers in the Compustat database, our sample drops to a quarter of the size. Despite having lower power, Column (3) presents a similar economic relationship between *Laboratory AI* and *Physical AI* compared to Columns (1) and (2), with F-statistics above 30 for all but one long-horizon specification.

IV.a Focal Firm Results

Table 6 reports the second stage results for the production by the non-IT focal firms as a function of whether the focal firm patents physical AI innovation. Across the columns are different dependent variables of production – log employment, log capital, revenue TFP, log output (equal to revenues minus change in inventory), and COGS/output. The rows, every three rows, represent individual estimations of cumulating linear projections. Because of the cumulation, we might expect the coefficients concavely to increase over the time horizons (because the years cumulate) until the largest economic magnitude effect, if any, is reached. We find this concave increasing to be largest and most robustly significant across columns for the two-year horizon, and hence we focus our interpretation largely on this horizon.

In the second stage, the coefficient on instrumented *PhysicalAI* is positive and statistically significant in explaining changes in log employment, log capital, TFP, and log output at the two-year horizon. After two years, the magnitudes of the coefficient on *PhysicalAI* diminish and become insignificant statistically. These estimates include firm fixed effects, commuting zone × year fixed effects, and SIC2 × year fixed effects, and control for firm size, idiosyncratic risk, lagged patenting, and the lagged level of the dependent variable.

It may seem unsurprising that innovation in physical AI is growth-oriented, since growth is the purpose of innovation. The mechanisms, however, are not obvious. For instance, innovation-driven growth might be accompanied by increases capital and labor, but physical AI innovation could just as well replace labor. Furthermore, physical AI innovation could raise profits through revenue productivity or cost efficiency.

Before we turn to economic magnitude effects from our estimation, we first consider Table 7, our incumbent analysis that removes firm-years in which inventors have recently moved to commuting zone. This is our preferred analysis in that the IV does not suffer from a selection of firms moving R&D activities to AI knowledge hubs. As before, we include firm fixed effects, commuting zone times year fixed effects, the log of firm recent patenting, and controls for firm size, risk, and the base level of the dependent variable.

Table 7 displays two clear differences from Table 6. First, the magnitudes of the coefficients on *PhysicalAI* at the two-year horizon are materially lower. Second, *PhysicalAI* no longer predicts an improvement in TFP. The fact that productivity gains are not robust across methods differentiates our findings from the automation literature, which emphasizes TFP growth through

capital augmentation and labor reallocation toward non-routine tasks (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, making use of the additional capital and labor we document.

Interpreting the coefficients in economic magnitudes, we find that a standard deviation (0.067 from Table 4 Panel A) higher *PhysicalAI* results in a 1.47% (= 0.437 x 0.067/2) higher employment growth rate (calculated as the raw average for the two-year cumulative projection divided by two) and a 3.14% higher capital growth rate. Note that we are reporting raw annual averages just to facilitate easy referencing, but the effects are increasing concave, diminishing over time, and subject to compounding effects in the growth rates.

The ratio of these figures is very much in line with the consensus from growth elasticity calculations from a Cobb-Douglas specification. To see this, we take the logarithmic derivative of the standard Cobb-Douglas equation of $Y = AK^{\alpha}L^{1-\alpha}$, yielding:

$$\frac{\Delta Y}{Y} = \frac{\Delta A}{A} + \alpha\frac{\Delta K}{K} + (1-\alpha)\frac{\Delta L}{L}. \quad (3)$$

The consensus from the long-standing literature has been that $\alpha \approx 0.33$ (Solow, 1957; Kaldor, 1961; Gollin, 2002). We can calculate the economic magnitude effect of the revenue results, both for interest in understanding growth generated by physical AI and for testing that the magnitude lines up with that which could be aligned with increased investment in capital and labor. In column (4) or Table 7, the coefficient on *PhysicalAI* as a predictor of output at two-year horizon is 0.626, significant at the 5% confidence level. Using the same calculation procedure, a one standard deviation increase in *PhysicalAI* results in 2.10% increase in the output growth rate. This semi-elasticity magnitude is in line with the growth in the input factors. Invoking equation (3) with $\Delta A = 0$ and $\alpha = 0.33$ inplies that $\frac{\Delta Y}{Y} = 0.33 * 3.14\% + .67 * 1.47\% = 2.02\%$. Our estimate of the change in the production growth rate as measured by revenues is 2.10%

The cost efficiency results in Table 6 and 7 were unambiguous in finding no evidence that physical AI innovation affects COGS/output in any horizon. Overall, our findings paint a picture of expansion rather than efficiency: physical AI innovation drives firms to scale up labor and capital in pursuit of revenue growth. This picture complements the findings of Babina et al. (2024).

Whereas their work brings out the importance of newly hired AI specialists, our evidence allows for the interpretation that physical AI innovations are implemented by domain-specific employees – such as engineers, production specialists, or technicians – who integrate AI tools into operational workflows. 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.

IV.b Downstream Customer Results

In this section, we step downstream from the focal firms to their customers. Customers could be affected by the upstream focal firm's physical AI innovation through different channels. The innovation can improve the quality of the products and processes the customer buys, such as when a hospital acquires the AI-enabled surgical device discussed in the introduction. The downstream benefits of physical AI innovation are not limited to medical devices. Consider a 2017 Deere & Company patent (US 9,719,973), which builds on prior AI patents on image analysis and pattern recognition, to compare images of a crop before and after a herbicide or fertilizer application and determines whether the application worked. A farmer who uses this system can learn which applications are effective, raise yield, and reduce wasted inputs.

In Table 8, we present the effect of upstream physical AI innovation on customer firms. As mentioned earlier, the sample size drops considerably because only large customers (accounting for more than 10% of sales) are disclosed and many disclosed customers are government organizations or private firms. As before, physical AI innovation is measured for the upstream focal firm and is instrumented on Laboratory AI of the focal firm. We control for the customer fixed effects, as well as the upstream firm commuting zone time year fixed effects, the customer SIC2 times year effects, and the customer log count of patents and firm size.

The customer results in Table 8 have a distinct feature not present in the focal firm results of Tables 6 and 7: there is a lag before customers show an impact from upstream physical AI. We find no discernible effects at horizons of one and two years. At three years, employment and output begin to rise, and at four years the estimates for employment, output, and TFP are all positive and significant, with cumulative effects peaking around year four for employment and output. A lag is plausible. Customers may be locked into existing contracts, processes, or products, and must first observe the upstream firm's change before adapting their own operations. Unlike for focal firms,

we also find that upstream physical AI predicts higher customer TFP. In the hospital example, this suggests that the device not only helps the hospital expand its surgical services and attract new patients, but may also help it make more efficient use of its existing labor and capital to generate revenue.

The economic magnitudes are muted relative to the focal firm results. Using the coefficients at the t+4 horizon, and the standard deviation of upstream physical AI from Panel B of Table 4 (0.057), a one standard deviation increase in *Upstream PhysicalAI* predicts a 0.40% higher annualized growth rate of employment, a 0.60% increased growth rate of capital, and a 0.45% increased TFP. If we again apply the Cobb-Douglas logarithmic derivative of equation (3), with $\alpha = 0.33$, $\frac{\Delta A}{A} = 0.45\%$ and filling in for the change in growth rate for capital and labor, we have $\frac{\Delta Y}{Y} = 0.45\% + 0.33 * 0.6\% + .67 * 0.4\% = 0.92\%$. Applying a one standard deviation change in Upstream PhysicalAI from Table 8 to the coefficient 0.575 and annualizing lead to an economic magnitude of 0.82% increase in output, which is in line with the 0.92% as predicted in the back-of-the-envelope Cobb-Douglas calculation.

IV.c Smart Automation Results

The literature on AI and labor markets has devoted substantial attention to automation.[25] From our vantage, it is important to note that automation is a broader concept than patented innovation: much of what firms implement as automation involves applying existing technologies to new settings rather than generating novel patented inventions. Moreover, automation patents explain little of the variance in physical AI (shown in Appendix Table A4). That said, automation and Physical AI are not orthogonal – approximately 35% of AI-citing patents outside the technology space meet the definition of automation patents.[26]

Table 9 decomposes physical AI innovation into its automation and non-automation components. Doing so allows us to examine the extent to which Physical AI innovation is directed towards automation and connect the emerging literature on AI to the seminal literature on routine automation. We calculate the fraction of a firm's patents that are classified both as Physical AI

[25] Acemoglu, Autor, Hazell, and Restrepo, 2022; Agrawal, Gans, and Goldfarb, 2025; Hampole, Papanikolaou, Schmidt, and Seegmiller, 2025; Brynjolfsson, Chandar, and Chen, 2025.

[26] 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.

patents and automation patents. We refer to this as $SmartAutomation$. Our choice of terminology reflects the fact that these patents embed AI capabilities – real-time pattern recognition, adaptive reasoning, dynamic response to variable conditions – that can extend automation into domains previously requiring human judgment.

We define *NonAutoPhysicalAI* as the fraction of a firm's patents classified as physical AI but not as automation patents. Separating the two speaks directly to the cost efficiency null in Tables 6 and 7: if any physical AI patents carry a cost-reduction effect, smart automation patents – where AI capabilities are directed at task automation – are the natural candidates. We explore whether $SmartAutomation$ operates through different production channels than other physical AI. Because we lack an instrument for the interacted specification of PhysicalAI with automation status, we limit this analysis to an OLS Linear Projection framework. Hence, we view the evidence as largely suggestive, provocative for future exploration.

Panel C of Table 9 reveals the central finding of this section: *SmartAutomation* strongly and robustly predicts cost efficiency, with a negative and significant coefficient on COGS/Output at both the three- and five-year horizons. The magnitude is economically meaningful. At a five-year horizon, a one-standard-deviation increase in automation physical AI predicts an annualized reduction in the COGS/output ratio of 0.0064.

To further unpack this result, we contrast smart automation with the routine automation emphasized in prior literature. A well-established literature argues that computers are adept at replacing labor in routine tasks that follow explicit, programmable rules (Autor, Levy, and Murnane, 2003). AI advances extend this automation frontier into tasks requiring cognitive judgment and real-time adaptation to changing conditions – precisely where *SmartAutomation* should generate cost efficiency gains. We therefore expect the cost efficiency effect to concentrate in industries with a high proportion of abstract and cognitive tasks. To test this, we use the industry-level Dictionary of Occupational Titles (DOT) task data of Autor, Levy, and Murnane (2003) and, following Autor, Katz, and Kearney (2006) and Autor and Dorn (2013), construct a variable *Abstract* that identifies industries with a high concentration of tasks requiring cognitive and abstract reasoning. Variable construction and the mapping to 3-digit SIC codes are described in Appendix C3.

Next, we assign firms to subsamples based on their primary 3-digit SIC code. Table 10, Panel A restricts the sample to firms in high-abstract/cognitive task industries; Panel B to firms in

low-abstract/cognitive task industries. The evidence confirms our hypothesis: the cost efficiency effect of *SmartAutomation* is concentrated in Panel A and absent in Panel B. That is, the cost-reducing effect of smart automation operates where traditional rule-based automation cannot – in industries where production requires judgment, pattern recognition, and real-time adaptation rather than the execution of predetermined routines.

Taken together, the results in Tables 9 and 10 qualify the broader physical AI narrative of our paper. Physical AI innovation generally drives revenue expansion rather than cost reduction. Smart automation carves out a meaningful exception: it is the subset of physical AI where AI absorption directly lowers the cost of output, and that effect is concentrated in sectors where cognitive and abstract reasoning are central to production. As AI capabilities continue to advance, this channel may grow in economic significance.

V. 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 physical AI innovation and measure it using citations by non-IT patents to prior AI patents. We then study physical AI both by the focal firms and in downstream dataset of customers of focal firms. Before doing so, we document the large importance of non-IT sectors, where most patenting happens. Our second motivating factoid is that U.S. firms appear to hold a consistently high advantage relative to European firms in physical AI innovation.

Our main analysis centers on the production inputs and outcomes of focal inventing firms and their customers. We find that firms engaging in physical AI innovation experience higher employment and capital growth, input of production. These augmented production factors turn into production revenues up to two years out. We find no conclusive evidence regarding the focal firm TFP and cost efficiency in the full sample. For customers, we find a lagged picture. Upstream physical AI predicts customers growth in production inputs and in TFP, amalgamating in approximately a 1% change higher output.

Finally, we show that smart automation — automation patenting that is physical AI innovation — has the opposite effect: increasing cost efficiency. These findings suggest that the economic value of physical AI innovation arises primarily through revenue expansion, except when the physical AI is smart automation. 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.

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.

Acharya, Viral V., Xuehua Shao, and Zhaoxia Xu. 2026. “AI and Productivity: The Role of Innovation.” Working Paper. New York University Stern School of Business, March.

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.

Alvarez, Fernando E., David Argente, Joyce Chow, and Diana Van Patten. 2026. “Data Centers and Local Economies in the Age of AI: A Shift–Share Approach.” Working Paper No. 35194. Working Paper Series. National Bureau of Economic Research, May.

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)*.

Angrist, Joshua D., and Jörn-Steffen Pischke. 2009. *Mostly Harmless Econometrics: An Empiricist’s Companion*. Princeton University Press.

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.

Arts, Sam, Lee Fleming, and Lena Veretennik. 2026. “Harnessing Academic Science for Corporate Technology: Interpersonal Networks and Absorptive Capacity.” *Organization Science* Forthcoming.

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, David Dorn, and Gordon Hanson. 2019. “When Work Disappears: Manufacturing Decline and the Falling Marriage Market Value of Young Men.” *American Economic Review: Insights* 1 (2): 161–78.

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 Melissa S. Kearney. 2006. “The Polarization of the U.S. Labor Market.” *American Economic Review* 96 (2): 189–94.

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.

Azoulay, Pierre, Joshua S. Graff Zivin, and Jialan Wang. 2010. “Superstar Extinction.” *The Quarterly Journal of Economics* 125 (2): 549–89.

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.

Balsmeier, Benjamin, Lee Fleming, and Sonja Lück. 2023. “Isolating Personal Knowledge Spillovers: Coinventor Deaths and Spatial Citation Differentials.” *American Economic Review: Insights* 5 (1): 21–34.

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., Erik Brynjolfsson, and Lorin M. Hitt. 2002. “Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence.” *The Quarterly Journal of Economics* 117 (1): 339–76.

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

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen. 2025a. “Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence.” *Working Paper*.

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. 2025b. “Generative AI at Work*.” *The Quarterly Journal of Economics* 140 (2): 889–942.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2021. “The Productivity J-Curve: How Intangibles Complement General Purpose Technologies.” *American Economic Journal: Macroeconomics* 13 (1): 333–72.

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.

Card, David. 2001. “Estimating the Return to Schooling: Progress on Some Persistent Econometric Problems.” Econometrica 69 (5): 1127–60.

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, Orlando Gomes, Roxana Mihet, and Kumar Rishabh. 2026. “Data Innovation Complementarity and Firm Growth.” *Management Science* Forthcoming.

Gollin, Douglas. 2002. Getting Income Shares Right. *Journal of Political Economy*, 110(2), 458–474.

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 B. Jaffe, and Manuel Trajtenberg. 2001. “The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools.” NBER Working Papers, NBER Working Papers, October, 8498.

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.

Jiang, Wei. 2017. “Have Instrumental Variables Brought Us Closer to the Truth?” The Review of Corporate Finance Studies 6 (2): 127–40.

Jordà, Òscar. 2005. “Estimation and Inference of Impulse Responses by Local Projections.” *American Economic Review* 95 (1): 161–82.

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.

Kaldor, Nicholas (1961). "Capital Accumulation and Economic Growth." In Friedrich A. Lutz & Douglas C. Hague (Eds.), The Theory of Capital (pp. 177–222). St. Martin's Press.

Karabarbounis, Loukas, & Neiman, Brent. 2013. The Global Decline of the Labor Share. *The Quarterly Journal of Economics*, 129(1), 61–103.

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.

Korinek, Anton, and Patrick McKelvey. 2026. "Measuring the AI Economy." Working Paper No. 26-9. Working Paper Series. Peterson Institute for International Economics, May.

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.

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.

Teece, David J. 1986. "Profiting from Technological Innovation: Implications for Integration, Collaboration, Licensing and Public Policy." *Research Policy* 15 (6): 285–305.

The Draghi Report on EU Competitiveness. (Luxembourg), European Union. 2024. https://commission.europa.eu/topics/competitiveness/draghi-report_en

Yue, Daniel, and Yiyang Zeng. 2026. "The Local Economic Effects of Data Center Entry." SSRN Scholarly Paper No. 6497238. Social Science Research Network, March 30.

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 1.B. depicts the total number of 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. See Appendix A for classification of IT patents. Figure 1.D. depicts the difference in the share of non-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.

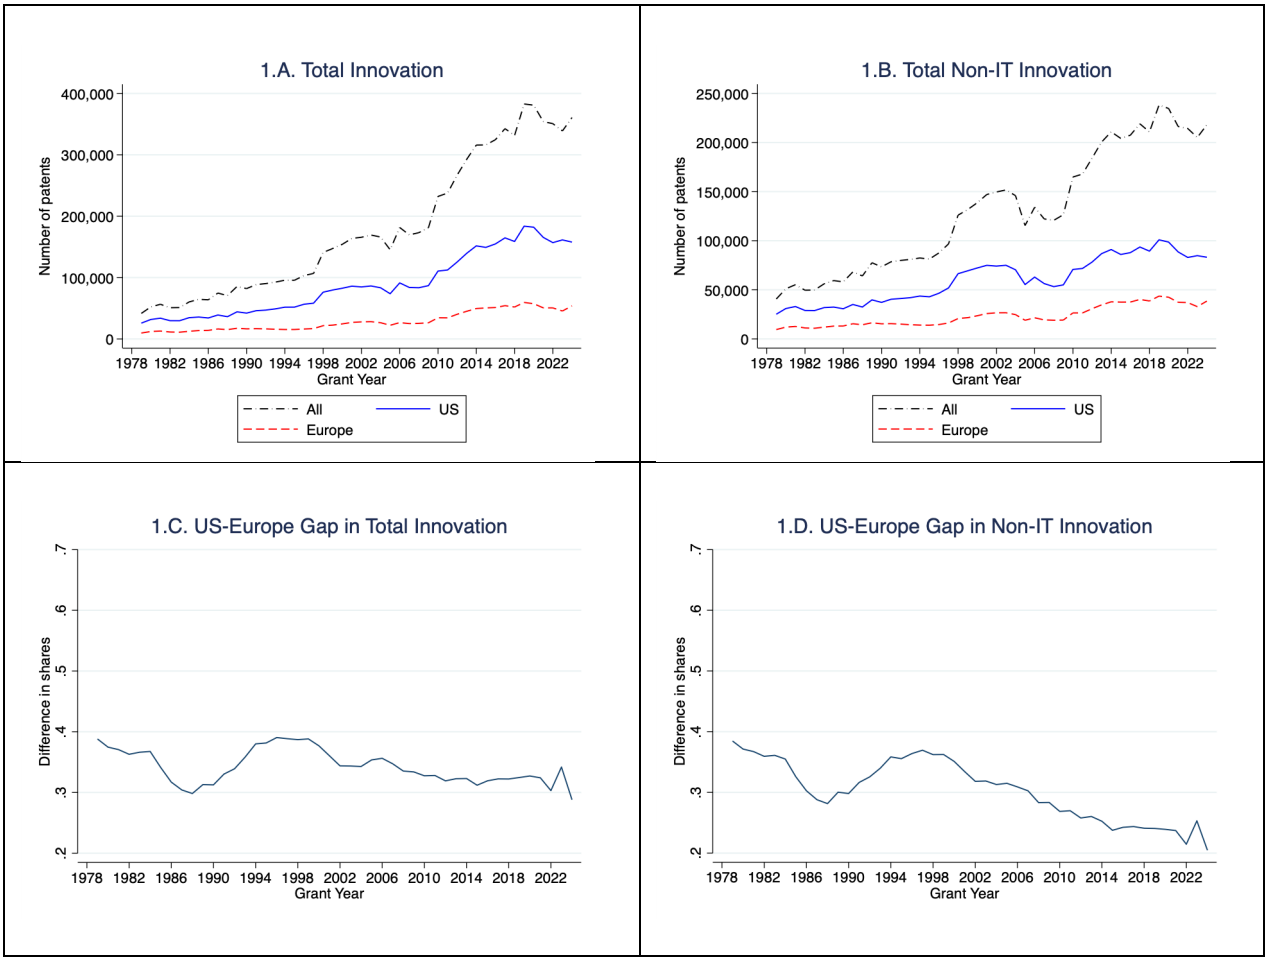


Figure 2: AI Innovation and Physical AI Innovation

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. Figure 2.C. depicts the difference in the share of total AI patents held by U.S. and European companies. Figure 2.B. depicts Physical AI innovation, defined as non-IT patents that cite AI patents. Figure 2.D. depicts the difference in the share of Physical AI innovation by U.S. and European companies. See Appendix A for classification of AI patents and IT patents.

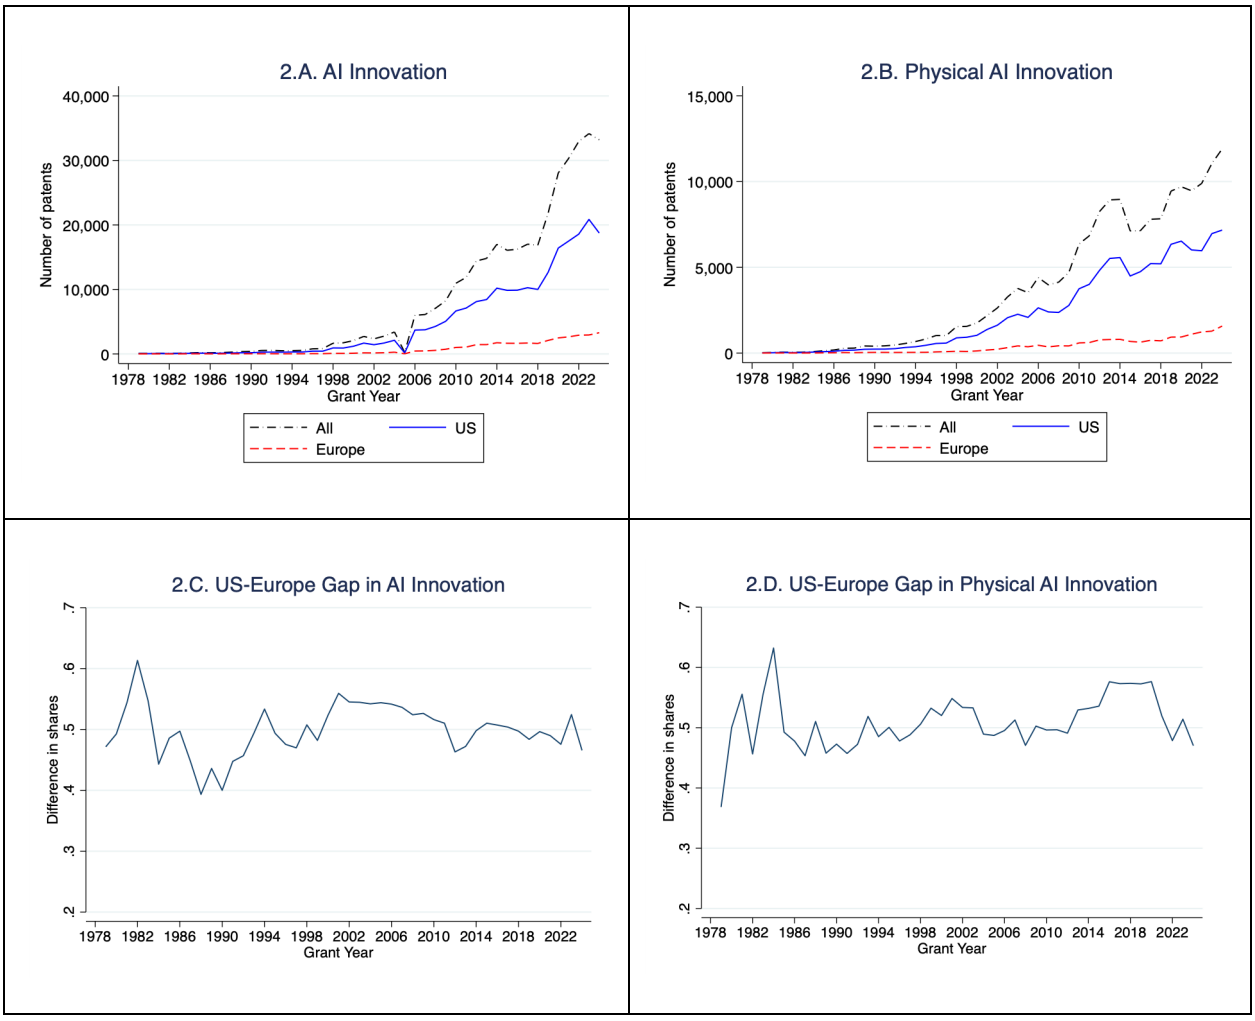


Figure 3: Heat Map of Laboratory AI Instrument

This map shows the log cumulative count of university and federal-lab patents that are either classified as AI patents or cite an AI patent, summed over the full sample period from 1979 to 2024 and aggregated to the 1990 Commuting Zone (CZ) level. CZ boundaries are constructed by mapping Census county boundaries to CZs using the county-to-CZ crosswalk provided by David Dorn (Autor, Dorn, and Hanson,2013), then spatially merging all counties sharing a CZ code into a single polygon. Darker shading indicates greater cumulative AI-related patent activity in universities and federal laboratories. We map the log-scale to make variation visible given the high geographic concentration of this activity.

The red circled commuting zones referenced in the paper are, from left to right, San Francisco Bay (SF, North Bay, East Bay); Iowa City/Davenport, the DMV region around Washington, DC.

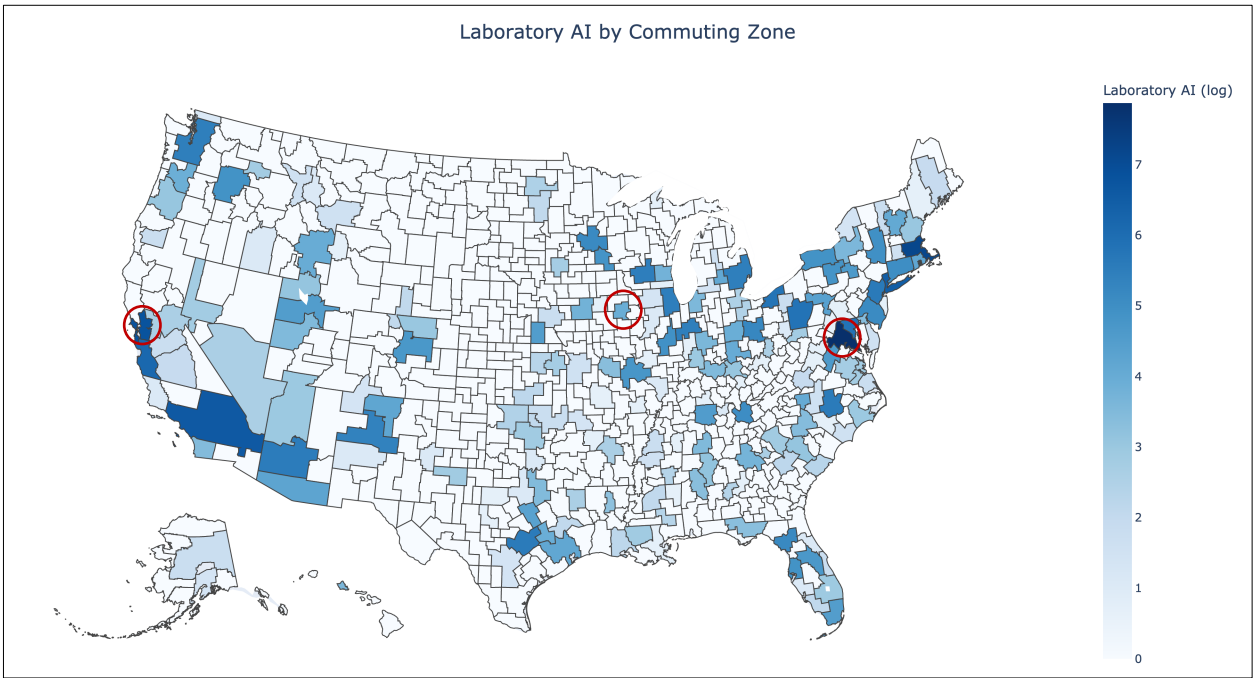


Table 1: Industry of Firms with Physical AI Innovation

Reported are the industries (2-digit SIC codes) of firms with physical AI innovation, i.e., firms that are not in the software and telecommunications sector and were granted at least one AI-citing patent between 1979 to 2024. *Freq.* is the number of unique COMPUSTAT firms within each 2-digit SIC with physical AI innovation, *Percent* is that count as a share of all firms with physical AI innovation, *Cum.* is the running cumulative percent with industries listed in descending order of frequency. SIC codes with at least 10 firms doing physical AI innovation are reported. The excluded software and telecommunications sector are firms with SIC codes are 357, 481, 482, 484, 489, and 737.

SIC2	SIC2 Description	Freq.	Percent	Cum.
38	Measuring, Analyzing, and Controlling Instruments; Photographic, Medical and Optical Goods	212	16.91	16.91
28	Chemicals and Allied Products	186	14.83	31.74
35	Industrial and Commercial Machinery and Computer Equipment	105	8.37	40.11
37	Transportation Equipment	80	6.38	46.49
36	Electronic and Other Electrical Equipment and Components	60	4.78	51.28
73	Business Services	50	3.99	55.26
87	Engineering, Accounting, Research, Management, and Related Services	42	3.35	58.61
60	Depository Institutions	33	2.63	61.24
99	Nonclassifiable Establishments	33	2.63	63.88
67	Holding and Other Investment Offices	31	2.47	66.35
49	Electric, Gas, and Sanitary Services	26	2.07	68.42
13	Oil and Gas Extraction	24	1.91	70.33
34	Fabricated Metal Products, Except Machinery and Transportation Equipment	23	1.83	72.17
39	Miscellaneous Manufacturing Industries	23	1.83	74
50	Durable Goods (Wholesale Trade)	23	1.83	75.84
62	Security and Commodity Brokers, Dealers, Exchanges, and Services	23	1.83	77.67
20	Food and Kindred Products	21	1.67	79.35
59	Miscellaneous Retail	21	1.67	81.02
63	Insurance Carriers	19	1.52	82.54
80	Health Services	18	1.44	83.97
27	Printing, Publishing, and Allied Industries	17	1.36	85.33
29	Petroleum Refining and Related Industries	16	1.28	86.6
26	Paper and Allied Products	15	1.2	87.8
61	Nondepository Credit Institutions	13	1.04	88.84
25	Furniture and Fixtures	11	0.88	89.71
30	Rubber and Miscellaneous Plastics Products	11	0.88	90.59
33	Primary Metal Industries	10	0.8	91.39

Table 2: 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. $\ln(\xi)$ is the Kogan et al (2017) patent value. 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. *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. See Appendix A for classification of AI patents and IT patents.

Panel A: All Patents

Variable	N	Mean	SD	Median	Min	Max
$\ln(\xi)$	2533690	0.977	2.263	1.445	-8.619	8.739
I_{AIcite}	2533690	0.161	0.367	0	0	1
ln (1+Forward)	2533690	0.415	0.545	0.240	0	7.145
ln Backward	2533690	1.988	1.162	1.946	0	8.630
AI Focal Patent	2533595	0.061	0.240	0	0	1
ln Market Cap	2345326	16.106	2.522	16.266	3.611	22.074
$\ln(\sigma)$	2343675	-4.087	0.467	-4.136	-5.501	0.922

Panel B: Non-IT Patents

Variable	N	Mean	SD	Median	Min	Max
$\ln(\xi)$	1,765,250	0.812	2.290	1.350	-8.619	8.510
I_{AIcite}	1,765,250	0.041	0.197	0	0	1
ln (1+Forward)	1,765,250	0.413	0.527	0.249	0	7.145
ln Backward	1,765,250	2.032	1.133	1.946	0	8.630
AI Focal Patent	1,765,250	0	0	0	0	0
ln Market Cap	1,621,447	15.623	2.443	15.793	3.611	22.074
$\ln(\sigma)$	1,620,311	-4.063	0.473	-4.112	-5.501	0.922

Table 3: Physical AI Innovation and Patent Value

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

$$ln(\xi_p) = \alpha + \beta\, I_{AIcite_p} + \gamma_1 \ln Market\ Cap_p + \gamma_2 \ln Backward_p + \mu_{i,t} + \rho_k + \epsilon_p$$

where ξ is the Kogan et al (2017) value of patent p. I_{AIcite} is an indicator variable taking the value 1 if the focal patent cites at least one AI patent. *ln Market Cap* is the firm's log market capitalization measured the day prior to patent grant date. *ln Backward* is the log of the number patents cited by the patent. $\mu_{i,t}$ are firm-by-grant-year fixed effects, ρ_k are technology class fixed effects. Column 1 (column 2) is restricted to the subsample of IT (non-IT) patents. The classification of IT patents is presented in Appendix .A ***, **,* indicate significance at the 1%, 5%, and 10% respectively. All variable definitions are provided in Appendix A. The sample period is 1979 to 2024.

	(1) IT Patents $ln(\xi)$	(2) Non-IT Patents $ln(\xi)$
I_{AIcite}	0.0014	0.0039**
	(1.32)	(2.15)
ln Market Cap	0.8679***	0.8900***
	(44.70)	(78.88)
ln Backward	0.0017***	0.0014***
	(3.18)	(3.31)
Observations	716,823	1,605,471
R-squared	0.97	0.96
Firm x Year FE	Yes	Yes
Patent Class Fe	Yes	Yes
Clustered SE	Firm x Yr	Firm x Yr

Table 4: Firm-level Summary Statistics

Panel A presents summary statistics for the sample of focal firms excluding the software and communications sector over the sample period 1979 to 2024. The level of observation is firm-year. 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 Compustat *sales* plus change in inventories (*invt*) deflated by the CPI. *COGS*/*Output* is calculated as Compustat *cogs*/(*sale*+change in inventory (*invt*)). 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. LaboratoryAI is the number of AI patents or AI-citing patents granted to universities and federal labs located in the firm's commuting zone.

Panel B presents the sample of customer firm-years. Customers of the focal firms in Panel A are identified from Compustat segment data. As per FASB Accounting Standards, this sample is restricted to customers accounting for more than 10% of focal firm sales and is conditional on a successful match of customer name to Compustat GVKEY. We drop customer firms in the software and communications industry and customer firm-years with positive Physical AI. *Physical AI Innovation (Focal firm)* is the percentage of the upstream firm's patents in a firm-year that cites an AI patent (same variable as Physical AI Innovation in Panel A). LaboratoryAI (Focal firm) is the number of AI patents or AI-citing patents granted to universities and federal labs located in the upstream firm's commuting zone (same variable as Laboratory AI in Panel A). All other variables in Panel B relate to customer firms.

Variable	N	Mean	SD	Median	Min	Max
Panel A: Focal Firm Sample						
Physical AI Innovation	263248	0.007	0.067	0	0	1
$I_{PhysicalAI}$	263248	0.030	0.170	0	0	1
ln Laboratory AI	263248	0.414	1.147	0	0	6.417
ln Employment	215887	-0.167	2.340	-0.124	-6.908	6.153
ln Capital Stock	220498	3.750	2.805	3.649	-6.999	11.211
Revenue TFP	145869	0.002	0.419	0.008	-7.961	5.552
ln Output	203349	4.654	2.442	4.723	-7.918	10.777
COGS/Output	204334	0.982	3.605	0.652	-0.242	99.735
ln (1+Patents)	240953	0.426	1.018	0	0	8.092
$\ln(\sigma)$	250538	-3.615	0.625	-3.649	-5.307	-1.673
Panel B: Customer Sample						
Physical AI Innovation (Focal firm)	44506	0.006	0.057	0	0	1
ln Laboratory AI (Focal firm)	44506	0.581	1.381	0	0	6.417
ln Employment	43910	3.531	1.859	3.871	-6.215	5.928
ln Capital Stock	44276	8.075	2.294	8.282	-6.304	11.921
Revenue TFP	38748	0.063	0.267	0.052	-6.563	3.725
ln Output	43167	8.793	1.915	9.127	-6.456	11.404
COGS/Output	43172	0.745	1.178	0.732	-0.014	62.732
ln (1+Patents)	43759	1.255	1.767	0	0	6.739
$\ln(\sigma)$	44327	-4.042	0.473	-4.101	-5.044	-2.033

Table 5: First Stage Results

Estimates of the following first stage regressions corresponding to 2SLS estimations shown in Tables 6, 7, and 8.

$$PhyscialAI_{it} = \alpha + \beta LaboratoryAI_{CZ(i),t} + \gamma X'_{it} + \delta_{jt} + \psi_{CZ,t} + \mu_i + \epsilon_{it}$$

Physical AI innovation is the fraction of the focal firm's non-IT patents that cite an AI patent. *LaboratoryAI* is the number of AI patents or AI-citing patents granted to universities and federal labs located in the firm's commuting zone. In all columns, $\psi_{CZ,t}$ are focal firm commuting zone times year fixed effects. In columns 1 and 2, μ_i are focal firm-fixed effects and δ_{jt} are focal firm SIC2 times year fixed effects. In column 3, μ_i are customer firm-fixed effects and δ_{jt} are customer SIC2 times year fixed effects. *X* is a vector of firm-time level control corresponding to the second-stage table indicated in the column title. t-statistics based on double-clustered standard errors are in parenthesis. ***, **,* indicate significance at the 1%, 5%, and 10% respectively. The sample period is 1979 to 2024.

	Dependent Variable:	Focal Firm Physical AI		Upstream Firm Physical AI
	Sample	Focal Firms		Customer Firms
Horizon		1 First stage for Table 6	2 First stage for Table 7	3 First Stage for Table 8
t+1	Laboratory AI	0.011*** (13.49)	0.012*** (12.62)	0.011*** (6.03)
	F-stat	181.87	159.24	36.33
	N	139093	132049	33498
t+2	Laboratory AI	0.011*** (12.55)	0.013*** (12.01)	0.011*** (5.86)
	F-stat	157.52	144.15	34.34
	N	125499	118747	31333
t+3	Laboratory AI	0.011*** (12.83)	0.012*** (12.20)	0.011*** (5.68)
	F-stat	164.69	148.72	32.31
	N	113660	107370	29490
t+4	Laboratory AI	0.011*** (12.14)	0.013*** (11.55)	0.010*** (5.55)
	F-stat	147.40	133.30	30.76
	N	103238	97385	27848
t+5	Laboratory AI	0.011*** (11.76)	0.012*** (10.84)	0.010*** (5.30)
	F-stat	138.25	117.59	28.10
	N	93881	88510	26288
	Controls	Yes	Yes	Yes
	Focal Firm FE	Yes	Yes	No
	Focal SIC2 x Year	Yes	Yes	No
	Focal CZ x Year	Yes	Yes	Yes
	Customer Firm FE	No	No	Yes
	Customer SIC2 x Year	No	No	Yes
	Clustering	Firm, Year	Firm, Year	Focal firm, Customer firm

Table 6: Physical AI Innovation and Focal Firm Outcomes

Estimates of the following instrumental variable estimation are presented

$$\ln(Y_{i,t+\tau}) - \ln(Y_{it}) = \alpha_0 + \beta_5 \widehat{PhysicalAI}_{it} + \gamma_1 X'_{it} + \delta_{jt} + \psi_{CZ,t} + \mu_i + \epsilon_{it}$$

where Y is the firm outcome listed in each column header. $\widehat{PhysicalAI}$ is the fitted value from the first stage regression shown in column 1 of Table 5. X is a vector of firm-time level control variables including log of dependent variable ($ln(Y_{i,t-1})$), log idiosyncratic volatility, log of firm size measured by current capital stock, and ln(1+Patents$_{i,t-1}$). δ_{jt} are SIC2 times year fixed effects and $\psi_{CZ,t}$ are commuting zone times year fixed effects. μ_i are firm-fixed effects. t-statistics based on standard errors double-clustered by firm and year are in parenthesis. ***, **,* 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: Δ in ...		Employment	Capital	TFP	Output	COGS/Output
Horizon of Δ (rows below):		(1)	(2)	(3)	(4)	(5)
t+1	$\widehat{PhysicalAI}$	0.584***	0.941***	0.146	0.490*	-0.660
		(4.62)	(4.48)	(0.60)	(1.80)	(-0.28)
	N	139093	140990	96393	123651	124828
t+2	$\widehat{PhysicalAI}$	0.542***	1.310***	0.697**	0.720**	-0.267
		(2.82)	(3.88)	(2.58)	(2.43)	(-0.11)
	N	125499	127311	86460	111755	112774
t+3	$\widehat{PhysicalAI}$	0.397	1.004**	0.541	0.247	-1.232
		(1.65)	(2.38)	(1.40)	(0.66)	(-0.47)
	N	113660	115391	78292	101657	102524
t+4	$\widehat{PhysicalAI}$	0.463	0.956*	0.229	0.207	-1.463
		(1.66)	(1.89)	(0.60)	(0.54)	(-0.93)
	N	103238	104780	71116	92566	93311
t+5	$\widehat{PhysicalAI}$	0.435	0.732	0.824*	0.368	-0.683
		(1.40)	(1.52)	(1.82)	(0.76)	(-0.44)
	N	93881	95319	64669	84365	85016
	Controls	Yes	Yes	Yes	Yes	Yes
	Firm FE	Yes	Yes	Yes	Yes	Yes
	SIC2 x Year	Yes	Yes	Yes	Yes	Yes
	CZ x Year	Yes	Yes	Yes	Yes	Yes
	Clustering	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year

Table 7: Physical AI Innovation and Focal Firm Outcomes: Incumbent Sample

Estimates of the following instrumental variable estimation are presented

$$\ln(Y_{i,t+\tau}) - \ln(Y_{it}) = \alpha_0 + \beta_5 \widehat{PhysicalAI}_{it} + \gamma_1 X'_{it} + \delta_{jt} + \psi_{CZ,t} + \mu_i + \epsilon_{it}$$

Where Y is the firm outcome listed in each column header. The sample excludes the first year *n* in which a firm has an inventor in a given commuting zone, as well as any year from *n+1* to *n+4* in which the firm has an inventor in that same commuting zone. $\widehat{PhysicalAI}$ is the fitted value from the first stage regression shown in column 2 of Table 5. X is a vector of firm-time level control variables including log of dependent variable ($ln(Y_{i,t-1})$), log idiosyncratic volatility, log of firm size measured by current capital stock, and ln(1+Patents$_{i,t-1}$). δ_{jt} are SIC2 times year fixed effects and $\psi_{CZ,t}$ are commuting zone times year fixed effects. μ_i are firm-fixed effects. t-statistics based on standard errors double-clustered by firm and year are in parenthesis. ***, **,* 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: Δ in ...		Employment	Capital	TFP	Output	COGS/Output
Horizon of Δ (rows below):		(1)	(2)	(3)	(4)	(5)
t+1	$\widehat{PhysicalAI}$	0.449***	0.740***	0.420*	0.590**	-2.047
		(3.17)	(3.52)	(1.91)	(2.50)	(-0.89)
	N	132049	133899	90585	117799	118921
t+2	$\widehat{PhysicalAI}$	0.437**	0.935***	0.434	0.626**	-0.595
		(2.50)	(2.97)	(1.33)	(2.02)	(-0.24)
	N	118747	120516	80966	106127	107090
t+3	$\widehat{PhysicalAI}$	0.301	0.887**	0.407	-0.032	0.270
		(1.20)	(2.09)	(0.99)	(-0.08)	(0.10)
	N	107370	109037	73148	96355	97166
t+4	$\widehat{PhysicalAI}$	0.327	0.877*	0.342	0.365	-1.305*
		(1.16)	(1.79)	(0.91)	(0.91)	(-1.78)
	N	97385	98892	66375	87626	88323
t+5	$\widehat{PhysicalAI}$	0.393	0.929*	0.755	0.291	1.129
		(1.10)	(1.74)	(1.62)	(0.57)	(0.46)
	N	88510	89902	60302	79773	80384
	Firm FE	Yes	Yes	Yes	Yes	Yes
	SIC2 x Year FE	Yes	Yes	Yes	Yes	Yes
	CZ x Year FE	Yes	Yes	Yes	Yes	Yes
	Clustering	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year

Table 8: Physical AI Innovation and Customer Firm Outcomes

Estimates of the following instrumental variable estimation are presented

$$\ln(Y_{i,t+\tau}) - \ln(Y_{it}) = \alpha_0 + \beta_5 \widehat{PhysicalAI}_{it} + \gamma_1 X'_{it} + \delta_{jt} + \psi_{CZ,t} + \mu_i + \epsilon_{it}$$

Y is the customer firm outcome listed in each column header. $\widehat{PhysicalAI}$ is the fitted value from the first stage regression shown in column 3 of Table 5. X is a vector of customer firm-time level control variables including log of dependent variable ($ln(Y_{i,t-1})$), log idiosyncratic volatility, log of firm size measured by current capital stock, and ln(1+Patents$_{i,t-1}$). δ_{jt} are customer SIC2 times year fixed effects, μ_i are customer firm-fixed effects, and $\psi_{CZ,t}$ are commuting zone times year fixed effects of the upstream focal firm. t-statistics based on standard errors double-clustered by customer firm and supplier firm are in parenthesis. ***, **,* 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: Δ in …		Employment	Capital	TFP	Output	COGS/Output
Horizon of Δ (rows below):		(1)	(2)	(3)	(4)	(5)
t+1	$\widehat{PhysicalAI}$	0.028	0.019	0.059	0.031	0.333
		(0.46)	(0.21)	(0.74)	(0.31)	(0.63)
	N	33498	33424	29253	32673	32702
t+2	$\widehat{PhysicalAI}$	0.047	0.123	0.124	0.042	0.756
		(0.74)	(1.14)	(1.54)	(0.42)	(0.80)
	N	31333	31245	27288	30547	30578
t+3	$\widehat{PhysicalAI}$	0.201**	0.200	0.074	0.317**	-0.475
		(2.08)	(1.41)	(0.70)	(2.07)	(-0.87)
	N	29490	29389	25715	28735	28767
t+4	$\widehat{PhysicalAI}$	0.280**	0.418*	0.314**	0.575**	-0.941
		(2.13)	(1.96)	(2.15)	(2.41)	(-1.17)
	N	27848	27747	24278	27162	27192
t+5	$\widehat{PhysicalAI}$	0.274**	0.233	0.373**	0.549**	0.284
		(2.01)	(1.07)	(2.14)	(2.29)	(0.22)
	N	26288	26189	22851	25645	25671
	Cust. Firm FE	Yes	Yes	Yes	Yes	Yes
	Cust. SIC2 x Year FE	Yes	Yes	Yes	Yes	Yes
	Focal CZ x Year FE	Yes	Yes	Yes	Yes	Yes
	Clustering	Cust. Firm, Focal Firm	Cust. Firm, Focal Firm	Cust. Firm, Focal Firm	Cust. Firm, Focal Firm	Cust. Firm, Focal Firm

Table 9: Smart Automation and Focal Firm Outcomes

This table presents the relation between *Smart Automation* 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:

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

where Y are firm outcomes represented in the Panel names; *SmartAutomation* 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.

	Horizon				
	t+1	t+2	t+3	t+4	t+5
	Panel A: Employment				
Smart Automation (Automation Physical AI)	0.027*	0.036	0.036	0.071	0.097*
	(1.79)	(1.21)	(0.93)	(1.41)	(1.70)
Non-Automation Physical AI	0.034**	0.056*	0.101***	0.105**	0.095
	(2.28)	(1.98)	(2.84)	(2.21)	(1.61)
Automation	0.021***	0.034***	0.039**	0.049**	0.045*
	(3.44)	(2.77)	(2.52)	(2.50)	(1.95)
Observations	164385	148238	134035	121597	110643
R-squared	0.057	0.073	0.083	0.092	0.100
	Panel B: Output				
Smart Automation (Automation Physical AI)	0.007	0.021	0.069	0.094	0.137**
	(0.33)	(0.67)	(1.51)	(1.63)	(2.03)
Non-Automation Physical AI	0.045	0.059	0.093	0.136**	0.069
	(1.62)	(1.26)	(1.65)	(2.17)	(0.88)
Automation	0.021*	0.034**	0.034	0.044*	0.059**
	(1.89)	(2.23)	(1.48)	(1.90)	(2.02)
Observations	146492	132119	120057	109327	99755
R-squared	0.075	0.101	0.106	0.113	0.121
	Panel C: COGS/Output				
Smart Automation (Automation Physical AI)	-0.012	-0.156	-0.192***	-0.168	-0.480**
	(-0.11)	(-1.37)	(-3.18)	(-0.91)	(-2.63)
Non-Automation Physical AI	0.092	-0.062	0.016	-0.382	0.035
	(0.37)	(-0.27)	(0.04)	(-1.61)	(0.11)
Automation	0.010	0.125**	0.164	0.179	0.198
	(0.10)	(2.03)	(1.45)	(1.34)	(1.64)
Observations	147888	133328	121099	110217	100535
R-squared	0.019	0.031	0.028	0.045	0.039
Industry x Year FE	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year

Table 10: Focal Firm Cost Efficiency & Smart Automation: Conditional on Abstract/Cognitive Tasks

This table subsamples the data by industry to learn more about the COGS/Output result in Table 9, panel C. Panel A (Panel B) includes only firms with primary SIC3 classified as having high (low) Abstract/Cognitive Task Intensity. Abstract/Cognitive Task Intensity is constructed using the Dictionary of Occupational Titles (DOT) task data of Autor, Levy, and Murnane (2003) and is described in Appendix C3. We present the coefficients β_τ for cumulative forward time periods $\tau \in \{1, \ldots, 5\}$ from estimations of the following model:

$$\frac{COGS_{i,t+\tau}}{Output_{i,t+\tau}} - \frac{COGS_{i,t}}{Output_{i,t}} = \alpha_0 + \beta_{1\tau} SmartAutomation_{it} + \beta_{2\tau} NonAutoPhysicalAI_{it} + \beta_{3\tau} Auto_{it} + \gamma_1 X'_{it} + \delta_{jt} + \epsilon_{it}$$

where $SmartAutomation$ 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 C1. 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 lagged COGS/Output. δ_{jt} are SIC1 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.

Y= COGS/Output	Horizon				
	t+1	t+2	t+3	t+4	t+5
	Panel A: Subsample Split: Industries with High Abstract/Cognitive Task Intensity				
Smart Automation (Automation Physical AI)	-0.195	-0.623**	-0.693**	-0.513	-1.315*
	(-0.37)	(-2.20)	(-2.36)	(-0.78)	(-1.96)
Non-Automation Physical AI	0.765	-0.090	0.811	-0.854**	-0.168
	(0.97)	(-0.24)	(0.71)	(-2.22)	(-0.37)
Automation	-0.089	0.339*	0.431	0.404	0.538
	(-0.27)	(1.95)	(1.19)	(0.96)	(1.46)
Observations	45119	40493	36593	33123	30028
R-squared	0.017	0.027	0.027	0.041	0.041
	Panel B: Subsample Split: Industries with Low Abstract/Cognitive Task Intensity				
Smart Automation (Automation Physical AI)	0.076	0.048	0.083	0.018	-0.098
	(0.78)	(0.45)	(1.30)	(0.16)	(-0.80)
Non-Automation Physical AI	-0.151	-0.078	-0.256	-0.178	0.041
	(-1.08)	(-0.53)	(-1.39)	(-1.50)	(0.11)
Automation	0.041	0.002	0.002	-0.011	-0.027
	(1.51)	(0.10)	(0.09)	(-0.53)	(-1.09)
Observations	90308	81845	74708	68353	62652
R-squared	0.012	0.031	0.037	0.080	0.032
Industry SIC1 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	Source	Description
Patent-level variables		
ξ	Kogan et al. data library	Kogan et al (2017) value of a patent
σ	CRSP	Idiosyncratic volatility in patent grant year calculated as in Kogan et al (2017)
AI Focal Patent	USPTO	Patent level indicator variable equal to 1 if focal patent is an AI patent. See Table A3 for classification of AI patents
Backward	USPTO	Patent level natural log of the number patents cited by the focal patent (i.e. backward citations)
Forward	USPTO	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
I_{AIcite}	USPTO	Patent level indicator variable equal to one if the patent cites at least one AI patent. See Table A3 for classification of AI patents.
$I_{Automation}$	USPTO	Patent-level 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.
$I_{Process}$	USPTO	Patent-level 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.
Market Cap	CRSP	Shares outstanding (shrout) times share price (prc) one day prior to patent grant date
Originality	USPTO	Calculated as one minus the HHI of the technology class distribution of the patent's backward citations, i.e., $1 - \sum_c s_{p,c}^2$ where $s_{p,c}$ is the share of patent p's backward citations in IPC subclass c.
Firm-level variables		
σ	CRSP	Idiosyncratic volatility in firm-year calculated as in Kogan et al (2017)
AI-Hiring	Babina et al. (2024)	Firm level Babina et al. (2024) measure of AI employment
Automation	USPTO	Firm-level fraction of a firm's patents that are automation patents. Appendix C.1. describes the classification of automation patents
Capital stock	COMPUSTAT	ppent, deflated by the NIPA price of equipment
COGS/Output	COMPUSTAT	cogs/(sale+change in inventory (invt))

Employment	COMPUSTAT	Number of employees (emp)
$I_{PhysicalAI}$	USPTO	Firm-level indicator variable equal to one for firm-years with at least one patent that cites an AI patent, and zero otherwise. See Table A3 for classification of AI patents.
LaboratoryAI	USPTO	For each inventor of patent *p* filed in year *t* we calculate local AI stock 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 from *t-5* to *t-1*. If a patent has multiple inventors, we use the highest value of local AI stock among its inventors. LaboratoryAI is the mean value of this local AI stock across all patents in a firm-year.
Output	COMPUSTAT	Calculated as Sales (sale) plus change in inventories (invt) and deflated by the CPI
Patents	USPTO	Number of patents filed in the firm-year
Process	USPTO	Firm-level fraction of a firm's patents with more than 50% process claims. Appendix C.2. describes the classification of process claims.
Profit	COMPUSTAT	Gross profit calculated as sales (sale) minus cost of goods sold (cogs) and deflated by CPI.
Revenue TFP	COMPUSTAT	Following Babina et. al (2024), Revenue TFP is estimated 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 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 Information Technology (IT) IPC Classes

IPC Class	Description
G06F	Electric digital data processing
G06K	Graphical data reading; Presentation of data; Record carriers; Handling record carriers
G06N	Computing arrangements based on specific computational models
G06Q	Information and communication technology (ICT) specially adapted for administrative, commercial, financial, managerial, supervisory or forecasting purposes
G06T	Image data processing or generation, in general
G06V	Image or video recognition or understanding
G10L	Speech analysis techniques or speech synthesis; Speech recognition
H04L	Transmission of digital information, e.g. telegraphic communication
H04M	Telephonic communication
H04W	Wireless communication networks

Table A3. Description of Artificial Intelligence (AI) IPC Classes

IPC Class	Description
G06F 18/	Pattern recognition
G06F 40/	Handling natural language data
G06K 9/	Methods or arrangements for recognizing patterns (pre 2013 IPC)
G06N 3/	Computing arrangements based on biological models
G06N 5/	Computing arrangements using knowledge-based models
G06N 7/	Computing arrangements based on specific mathematical models
G06N 10/	Quantum computing, i.e. information processing based on quantum-mechanical phenomena
G06N 20/	Machine learning
G06T 7/	Image analysis
G06V	Image or video recognition or understanding
G10L 13/027	Concept to speech synthesis; Generation of natural phrases from machine-based concepts
G10L 15/	Speech recognition
G10L 17/	Speaker identification or recognition
G10L 25/63	Detecting or assessing emotion or mental state from voice signals
G10L 25/66	Detecting or assessing mood or disposition from voice signals
G06F 15/16, /76, /80, /82, /163, /167, /173, /177	Architectures for parallel and distributed digital computing, including multi-processor configurations, interconnection networks, and array processing units

Table A4: Is Physical AI a Proxy for Automation, Process Patents, or AI Hiring?

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	Median	Min	Max
Patent level						
I_{AIcite}	1,765,250	0.041	0.197	0	0	1
$I_{Automation}$	1,763,640	0.138	0.344	0	0	1
$I_{Process}$	1,765,237	0.284	0.451	0	0	1
Firm level						
Physical AI Innov.	263,248	0.0073	0.0671	0	0	1
Automation	263,248	0.0353	0.1488	0	0	1
Process	263,248	0.0640	0.2014	0	0	1
ln AI-Hiring	40,548	0.7130	1.2847	0	0	9.429

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})			
$I_{Automation}$	0.075***	0.038***		
	(14.13)	(14.48)		
$I_{Process}$			0.015***	0.014***
			(9.32)	(14.41)
Observations	1,763,734	1,746,615	1,765,332	1,748,218
Partial R-sq	0.017	0.005	0.001	0.001
Fixed effects	No	Firm x Year, Class	No	Firm x Year, Class
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	0.214***	0.191***				
	(0.014)	(0.015)				
Process			0.087***	0.080***		
			(0.008)	(0.008)		
ln AI-Hiring					0.022***	0.020***
					(0.002)	(0.003)
Observations	238197	233528	238197	233528	39347	39300
Partial R-sq	0.118	0.107	0.029	0.027	0.033	0.023
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 A5: Relation between Laboratory AI and Firm Outcomes in Placebo Sample

Coefficients β_τ from this reduced form regression are presented for cumulative forward time periods $\tau \in \{1, \ldots, 5\}$

$$\ln(Y_{i,t+\tau}) - \ln(Y_{i,t}) = \alpha_0 + \beta_\tau LaboratoryAI_{i,t} + \gamma_1 X'_{i,t} + \delta_{j,t} + \psi_{CZ,t} + \epsilon_{i,t}$$

The sample is restricted to firms with at least one patent, but zero AI-citing patents. Y are firm outcomes listed in the title of each panel. *LaboratoryAI* is the instrument used in our LP-IV analysis and is described in Appendix Table A1. X represents the same set of control variables used in main specifications shown in Tables 5 through 10. $\delta_{j,t}$ are SIC2 times year fixed effects and $\psi_{CZ,t}$ 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: Employment		
Laboratory AI	-0.000	0.000	0.001	0.001	-0.001
	(-0.03)	(0.10)	(0.16)	(0.13)	(-0.13)
Observations	28330	26745	24817	22871	21029
R-squared	0.190	0.206	0.221	0.233	0.236
			Panel B: Capital		
Laboratory AI	0.001	0.002	0.002	0.002	-0.000
	(0.56)	(0.52)	(0.27)	(0.20)	(-0.02)
Observations	28510	26940	25008	23036	21177
R-squared	0.198	0.221	0.244	0.260	0.265
			Panel C: TFP		
Laboratory AI	0.001	0.004	0.004	-0.001	0.006
	(0.35)	(0.76)	(0.58)	(-0.18)	(0.90)
Observations	24259	22592	20898	19249	17680
R-squared	0.166	0.209	0.241	0.297	0.316
			Panel D: Profit		
Laboratory AI	0.002	0.006	0.003	0.007	0.007
	(0.64)	(1.15)	(0.55)	(1.18)	(0.92)
Observations	23847	22526	21019	19461	17963
R-squared	0.239	0.259	0.278	0.282	0.293
			Panel E: Output		
Laboratory AI	-0.001	-0.005	-0.005	-0.011	-0.008
	(-0.15)	(-1.13)	(-0.72)	(-1.42)	(-1.03)
Observations	25583	24075	22367	20695	19090
R-squared	0.171	0.209	0.218	0.232	0.241
			Panel F: Cost Efficiency		
Laboratory AI	0.004	0.033	0.022	0.038	0.067
	(0.10)	(1.10)	(0.34)	(0.78)	(1.01)
Observations	25764	24249	22523	20822	19187
R-squared	0.116	0.124	0.141	0.150	0.148
Industry x Year FE	Yes	Yes	Yes	Yes	Yes
CZ x Year FE	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year

Table A6: Physical AI Innovation and Focal Firm Outcomes: Dropping AI Hubs

Estimates of the following instrumental variable estimation are presented after dropping five commuting zones that are have the highest intensity of Laboratory AI: San Francisco–Oakland, San Jose/Silicon Valley, Austin, Boston, and the Greater DC/Bethesda/Northern Virginia area. In addition, the sample excludes the first year *n* in which a firm has an inventor in a given commuting zone, as well as any year from *n+1* to *n+4* in which the firm has an inventor in that same commuting zone *n*.

$$\ln(Y_{i,t+\tau}) - \ln(Y_{it}) = \alpha_0 + \beta_5 \widehat{PhysicalAI}_{it} + \gamma_1 X'_{it} + \delta_{jt} + \psi_{CZ,t} + \mu_i + \epsilon_{it}$$

where Y is the firm outcome listed in each column header. $\widehat{PhysicalAI}$ is the fitted value from the first stage regression shown in column 1 of Table 5. X is a vector of firm-time level control variables including log of dependent variable ($ln(Y_{i,t-1})$), log idiosyncratic volatility, log of firm size measured by current capital stock, and ln(1+Patents$_{i,t-1}$). δ_{jt} are SIC2 times year fixed effects and $\psi_{CZ,t}$ are commuting zone times year fixed effects. μ_i are firm-fixed effects. t-statistics based on standard errors double-clustered by firm and year are in parenthesis. ***, **,* 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: Δ in ...		Employment	Capital	TFP	Output	COGS/Output
Horizon of Δ (rows below):		(1)	(2)	(3)	(4)	(5)
t+1	$\widehat{PhysicalAI}$	0.881***	1.088***	0.421	1.143***	-4.158
		(3.90)	(2.98)	(0.96)	(2.72)	(-1.02)
	N	110007	111752	71357	97286	98319
t+2	$\widehat{PhysicalAI}$					
		0.898**	1.206**	0.874	1.172*	3.542
	N	(2.36)	(2.02)	(1.45)	(1.99)	(0.82)
t+3	$\widehat{PhysicalAI}$	97885	99549	62981	86765	87645
	N	0.680	1.559**	0.333	0.874	-4.692
t+4	$\widehat{PhysicalAI}$	(1.31)	(2.20)	(0.51)	(1.25)	(-1.01)
		87874	89437	56367	78193	78925
	N					
t+5	$\widehat{PhysicalAI}$	0.822	2.137**	0.298	1.082	-3.636
		(1.32)	(2.43)	(0.54)	(1.54)	(-0.99)
	N	79256	80648	50753	70727	71362
	Controls	Yes	Yes	Yes	Yes	Yes
	Firm FE	Yes	Yes	Yes	Yes	Yes
	SIC2 x Year	Yes	Yes	Yes	Yes	Yes
	CZ x Year	Yes	Yes	Yes	Yes	Yes
	Clustering	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year

Table A7: Physical AI Innovation and Focal Firm Outcomes: Innovative firms only

Estimates of the following instrumental variable estimation are presented focusing only on the sample of firms that have at least one patent over the sample period.

$$\ln(Y_{i,t+\tau}) - \ln(Y_{it}) = \alpha_0 + \beta_5 \widehat{PhysicalAI}_{it} + \gamma_1 X'_{it} + \delta_{jt} + \psi_{CZ,t} + \mu_i + \epsilon_{it}$$

where Y is the firm outcome listed in each column header. $\widehat{PhysicalAI}$ is the fitted value from the first stage regression shown in column 1 of Table 5. X is a vector of firm-time level control variables including log of dependent variable ($ln(Y_{i,t-1})$), log idiosyncratic volatility, log of firm size measured by current capital stock, and ln(1+Patents$_{i,t-1}$). δ_{jt} are SIC2 times year fixed effects and $\psi_{CZ,t}$ are commuting zone times year fixed effects. μ_i are firm-fixed effects. t-statistics based on standard errors double-clustered by firm and year are in parenthesis. ***, **,* 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: Δ in ...		Employment	Capital	TFP	Output	COGS/Output
Horizon of Δ (rows below):		(1)	(2)	(3)	(4)	(5)
t+1	$\widehat{PhysicalAI}$	0.538***	0.921***	0.112	0.426	-0.526
		(4.36)	(4.35)	(0.43)	(1.46)	(-0.20)
	N	72597	73146	58916	66300	66774
t+2	$\widehat{PhysicalAI}$					
		0.444**	1.263***	0.659**	0.615*	-0.209
	N	(2.22)	(3.52)	(2.22)	(1.91)	(-0.08)
t+3	$\widehat{PhysicalAI}$	67077	67630	53958	61192	61611
	N	0.271	0.955**	0.466	0.155	-1.102
t+4	$\widehat{PhysicalAI}$	(1.09)	(2.24)	(1.17)	(0.40)	(-0.40)
		61985	62550	49726	56680	57059
	N					
t+5	$\widehat{PhysicalAI}$	0.431	0.982*	0.087	0.144	-1.708
		(1.48)	(1.89)	(0.22)	(0.36)	(-1.01)
	N	57322	57850	45855	52508	52823
	Controls	Yes	Yes	Yes	Yes	Yes
	Firm FE	Yes	Yes	Yes	Yes	Yes
	SIC2 x Year	Yes	Yes	Yes	Yes	Yes
	CZ x Year	Yes	Yes	Yes	Yes	Yes
	Clustering	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year

APPENDIX B

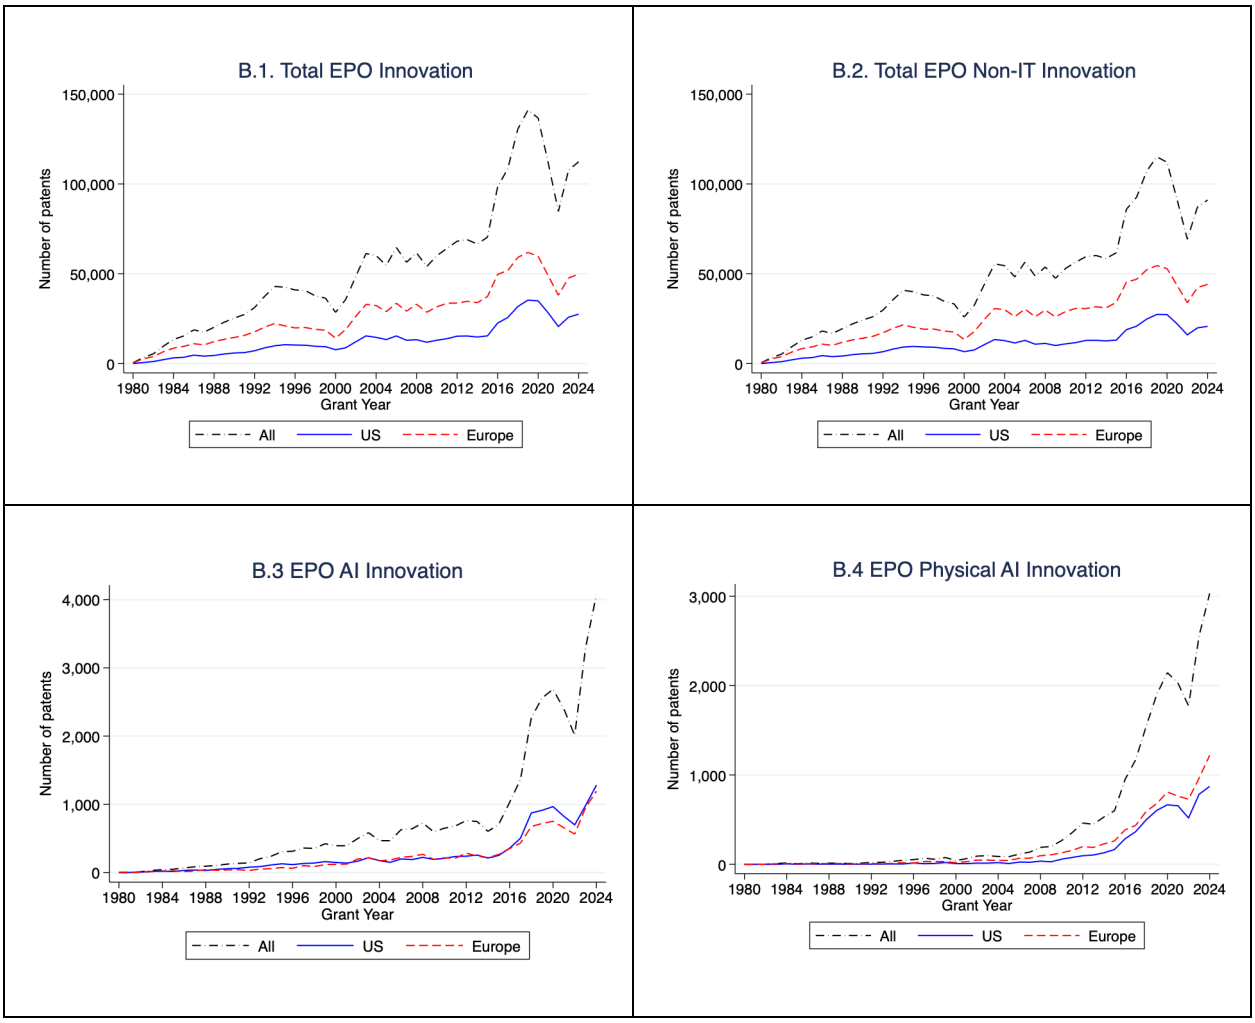
Data From European Patent Office

In Section IIa, we demonstrate that within the USPTO data, the U.S. leads Europe in physical AI innovation by 50 percentage points suggesting that U.S. firms are outperforming their European counterparts in integrating AI into the broader economy. To address concerns that U.S. dominance in physical AI 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). Specifically, the European advantage in EPO patents is between 20 to 30 percentage points in both overall innovation and non-IT innovation.

If U.S. dominance in AI innovation and physical AI innovation within the USPTO data (Figure 2) were primarily an artifact of this home-market bias, one would expect European firms to exhibit a parallel AI 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 physical AI patents issued by the EPO. Both figures show that U.S. firms are broadly comparable to European firms. The fact that European firms fail to command a lead in physical AI 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. 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. Figure B.4 depicts Physical AI EPO patents, defined as non-IT patents that cite AI patents. See Appendix A Tables A2 and A3 for the classification of IT patents and AI patents.



Appendix C: Automation Patents, Process Patents, and Cognitive/Abstract Occupations

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.

C3. Classification of Cognitive/Abstract Occupation

Our classification of abstract occupations is based on the *Dictionary of Occupational Titles* (DOT) task constructs provided by Autor, Levy and Murnane (2003). Specifically, we use their two DOT variables that measure non-routine cognitive tasks. The first, DCP (Direction, Control, and Planning) captures activities that are interactive and require managerial and communication skills. The second, GED-Math, measures quantitative reasoning requirements. Details of these variables are provided in Appendix 1 of Autor, Levy and Murnane (2003). Next, following Autor, Katz and Kearny (2006) and Autor and Dorn (2013), we create the abstract task measure, *Abstract*, as the simple average of DCP and GED-Math. Industries with above (below) -median values of *Abstract* are classified as having High (Low) Abstract/Cognitive Task Intensity

Autor et al (2003) provide industry-level task measures using the Census Industry Classification (CIC). We link this industry scheme to 4-digit SIC (1987 revision) through crosswalks provided by Autor, Dorn, and Hanson (2019). Because Compustat frequently reports 3-digit "group" SIC codes with no corresponding 4-digit entry in the underlying crosswalk, we aggregate the resulting industry-level task measures to the 3-digit SIC level, averaging across all 4-digit sub-industries within each 3-digit group; this substantially improves match coverage relative to matching at the 4-digit level. The resulting measure uses task scores from the 1990 Census cross-section — the most recent decennial vintage available in the Autor et al data — and

is merged onto the firm-year sample by 3-digit SIC code as a time-invariant industry characteristics.