Physical AI Innovation*

Adair Morse
University of California - Berkeley
& NBER

Amrita Nain
University of Iowa

August 2026

Abstract

We examine *physical* AI — the incorporation of artificial intelligence (AI) advances into inventions across the non-technology economy. Physical AI patenting is nearly as prevalent as technology AI patenting, commands excess financial market value and reflects 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 0.16%-1.02% and 0.17%-2.13%, respectively. Increases of similar magnitude are observed in profit growth, driven primarily through a revenue growth channel. One important exception arises: physical AI patents that are also automation patents – which we term *smart automation* – reduce 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. This paper was previously titled AI-Absorbing Innovation.

I. Introduction

Since 2005, the S&P 500 has delivered an annualized return of 10.7%, compared to 7.8% for the STOXX Europe 600.[1] Much of this differential has been attributed to the United States's continuing advantage in producing frontier information technology (IT).[2] However, unnoticed in this attribution is the U.S. advantage in physical AI innovation, the incorporation of advances in artificial intelligence (AI) into innovation outside the IT sector. Eighty percent of patenting occurs outside IT patent classes, and physical AI patenting is comparable in frequency to AI innovation itself. Yet, with a few notable exceptions, the literature on the economic effects of AI has developed in isolation from how AI functions as a general-purpose technology being absorbed into innovation across the broad real economy where most patenting occurs.

We show that physical AI innovation is both prevalent and economically valuable, and that U.S. firms lead Europe by a 50-percentage point margin. We study the production implications of physical AI innovation – how it affects the factors of production (labor, capital), production outputs (profits, revenues), total factor productivity (TFP), and cost efficiency. Our findings suggest that the value it generates operates primarily through output expansion rather than efficiency gains, a distinction with direct implications for how the AI revolution will shape employment, and economic growth.

Our work builds on two strands of research: the firm-level literature examining AI's economic effects through labor substitution, task augmentation, and productivity,[3] 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, Mihet, Rishabh, and Gomes (2024), and Argente, Baslandze, Hanley, and Moreira (2025), each of which bridges the two literatures by examining how AI shapes 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.

---

[1] As of March 2026

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

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

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

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

The example makes 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. Further, physical AI is broader than the common misconception 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, systems in agricultural, energy, and infrastructure, among many others.

Before turning to production implications, we characterize physical AI along two dimensions. First, using the Kogan, Papanikolaou, Seru, and Stoffman (2017) framework, we find that physical AI patents command a 40-basis-point valuation premium relative to comparable patents within the same firm-year and technology class. Together with the large and persistent transatlantic gap documented in our figures, this suggests that AI absorption into physical innovation is a key driver of cross-country growth differentials and may account for a portion of the sustained outperformance of U.S. equity markets.

The second important upfront characterization addresses whether our physical AI measure simply re-labels 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 Babina et al. (2024)'s measure of AI-skilled hiring is also low. This is not

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

surprising: hiring is an input measure, while physical AI – as captured through patent citations – is output evidence that AI knowledge has been incorporated into non-AI innovation.[6] Empirically, these could be two sides of the same coin, but they are not. 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.

Our main analyses examine how physical AI innovation alters the firm's production frontier. We trace its effects through real-economy inputs (labor and capital) and production outputs (profitability and revenues), while explicitly decoupling these outcomes from changes in productivity and cost efficiency. In doing so, we contribute to the growing body of research studying AI's consequences for firm growth, productivity, and asset dynamics.[8]

We identify physical AI innovation through patent citations, specifically non-IT patents that incorporate AI advances, as measured by their citations to AI patents. This approach follows a large literature that interprets patent citations as proxies for knowledge spillovers and technological knowledge flows (e.g., Jaffe, 1986; Jaffe, Trajtenberg, and Henderson, 1993; Hall, Jaffe, and Trajtenberg, 2001). Our analysis uses 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. In the subset of patenting firms in the non-IT sector, 7% have at least one physical AI patent.

Our first set of empirical analyses examines the factors of production (labor and capital) and factor productivity (revenue TFP). Our empirical methodologies follow the local projection regression (LP) models of Jorda (2005) and an instrumental variables implementation of linear projections (LP-IV) following Ramey (2016). Without instrumenting, the LP results may reflect firm selection on the availability of profitable innovation opportunities and on their stock of

---

[6] Our measurement is similarly tied to Fedyk et al. (2024), who study a particular instance of the AI-innovation nexus through human capital investment via data security specialists' role in non-security R&D teams.

[7] Corporate engineers (not so different from research economists) may learn AI skills sufficiently to deploy prior AI advancements in their non-AI innovation. Alternatively, akin to standard corporate practice to contract litigation lawyers rather than hire such skills in-house for one-off needs, the R&D departments may simply contract expertise to incorporate AI expertise in innovation.

[8] E.g., 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., 2024; and Baslandze et al. (2026).

knowledge capital, which enables them to recognize and successfully exploit the potential value of AI-using physical innovation.

Our LP-IV strategy exploits 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. Conditional on industry and region controls, the local supply of AI knowledge from academic and government research institutions is plausibly unrelated to firm-specific shocks affecting outcomes, satisfying the exclusion restriction. More than 40% of patents involve inventors located in a different commuting zone than the firm's headquarters, creating within–commuting-zone variation in Laboratory AI across firms. This allows us to include commuting zone times year fixed effects, which absorb time varying location effects that could otherwise compromise the exclusion restriction.

Our production input results reveal consistent positive effects of physical AI on the factors of production. A one-standard-deviation increase in physical AI innovation implies annualized employment growth of 0.16% to 1.02% over a five-year horizon, and annualized capital stock growth of 0.17% to 2.13% over a three-year horizon. The capital effect becomes statistically insignificant beyond three years, while the employment effect persists through the full five-year horizon. TFP effects are absent in the LP specification but emerge in the LP-IV, with significant positive effects at multiple horizons.

Turning to the outcomes from production, we first look at profits, where we find that physical AI results in increased profit levels over the full five-year horizon. A one standard deviation increase in physical AI innovation results in profit growth increases of 0.22%-1.80% per annum over the horizon.

Taken together and interpreted under the assumptions of our identification strategy, physical AI innovation induces firms to scale up labor and capital in pursuit of new opportunities, generating sustained profit growth. The annualized magnitudes are economically meaningful and internally consistent across outcomes - the growth effects on labor, capital, and profits are of similar scale across outcomes and specifications. We note that the LP-IV estimates are larger, consistent with a local average treatment effect (LATE) interpretation— the instrument identifies firms that actively seek out AI knowledge. Jiang (2017) suggests in a similar setting that the instrument may treat those for whom the effects are likely strongest. As Jiang (2017) conveys, the instrument is still valid under the exclusion restriction, yet the LATE effect may reflect the higher portions of a heterogenous treatment effects distribution. Hence, we take the range (from the LP to the LP-IV) as indicative of the distribution of effects that would be natural across firms.

With the production factor and profit results established, we turn to the mechanisms underlying profit growth. Does physical AI drive revenue expansion or cost reduction?

We find that physical AI predicts output growth, measured as sales plus the change in inventory, but shows no significant effect on COGS/Output. Hence, our evidence suggests that firms invest in physical AI innovation because of expansion opportunities. 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 – factor accumulation, profitability, and the cost structure – using a quasi-experimental identification strategy.

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.[9]

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

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.[10] A first take-away from these simple graphs is the importance of the non-IT sector in overall innovation. Over time, the 6.7 million non-IT patents represent almost 80% of the total. While media attention is concentrated on innovation in the IT sector, the non-IT sector accounts for the majority of patenting activity in the U.S. economy. Therefore, understanding how these sectors absorb and translate AI knowledge into value and employment creation is of 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 to motivate a second contribution of our work. Since 2005, the S&P 500 has delivered an annualized return of 10.7%, compared to 7.8% for the STOXX 600 (as of May, 2026), a gap widely attributed to the United States' superior innovative performance. Understanding where innovation-driven value originates speaks directly

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

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

to the U.S.–Europe performance gap and, from the European perspective, to potential remedies. Our focus on non-IT sectors is particularly germane as European firms hold comparative advantages in manufacturing, industrials, and other physical-economy sectors – domains where physical AI innovation could prove consequential.

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. Our second graphical punchline is that Europe has been closing the gap in non-IT innovation.[11]

In Figures 2.A and 2.B, we plot total AI patents granted and total AI-absorbing patents, respectively.[12] Following the citation-based convention for measuring knowledge absorption, we define AI-absorbing patents (or, interchangeably, physical AI innovation) as non-IT patents that cite a prior AI patent.

Physical AI innovation, as we use the term, does not refer to implementing an off-the-shelf AI-driven robotic or autonomous systems. Rather, it refers to patenting, the innovation, in physical-economy sectors, such as manufacturing, aerospace, energy, agriculture, etc., that build in AI technologies, as evidenced by citation to a prior AI patent. Because the commercial value of such inventions accrues in the application domain, patent classifications follow accordingly: 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

[11] One concern is that this pattern could reflect European firms increasingly filing non-IT patents at the USPTO over time rather than actual convergence in innovative activity; we address this in Appendix B using European Patent Office (EPO) data.
[12] See appendix Table A3 for the classification of AI patents. Robustness to a narrower definition of AI patents is presented in Appendix D.

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.[13] 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 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

[13] 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.

automation,[14] process patenting,[15] or AI hiring.[16] 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 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.[17]

---

[14] 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).

[15] 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.

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

[17] 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.

II.d Characterizing Physical AI Innovation: Market Value

To further characterize physical AI innovation, we compare AI-citing patents to other patents in terms of creative and economic value. We use two measures: patent originality (creative value) calculated following Hall, Jaffe, and Trajtenberg (2001) and the patent market value (economic value), measured by the (ξ) of Kogan et al. (2017). *Originality* is 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. Thus, patents whose backward citations belong to a narrow set of technologies have lower originality scores.

We create 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 $\ln(\xi)$, forward and backward citations, firm size, volatility are similar across Panel A and B.

We estimate the relation between AI absorption and value of a mean patent *p* granted to firm *i* as follows:

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

Where *V* is either the patent's creative value as measured by *Originality* or its log market value measured by $ln(\xi_p)$. $ln\, Backward$ is the log number of patents cited by the patent, included because assimilation of any existing knowledge, including non-AI 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 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.

Columns 1 and 2 of Table 3 show estimates of equation (1) in the subsample of IT patents only, while columns 3 and 4 focus on the subsample of non-IT patents. In both subsamples, I_{AIcite} has a positive and significant coefficient when the dependent variable is Originality. That is, patents that absorb AI technology have higher creative value both within the IT and non-IT patent classes. This finding is consistent with AI being a rapidly evolving technology: any patent that draws on AI knowledge - whether an IT or non-IT patent - is engaging with a novel and technically diverse knowledge base.

However, when the dependent variable is $ln(\xi_p)$, the coefficient on I_{AIcite} is insignificant in the IT subsample. That is, AI absorbing IT patents do not have higher market valuations than other IT patents (column (2)). In contrast, within the sample of non-IT patents, AI absorbing patents (i.e., physical AI citations) command higher economic value than other non-IT patents, significant at the 5% level. The inference from the β coefficient in column (4) 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.[18] Our finding is an overlooked inference: 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.

Together the findings in Table 3 and Figure 2.D point to AI absorption outside the technology space being value-enhancing, with U.S. firms enjoying a large transatlantic gap in such innovation. The combination of a valuation premium on physical AI patents and a persistent U.S. lead in producing them points to AI absorption outside the technology sector as a potential driver of the transatlantic market valuation gap.

II.e Firm-level measures of Physical AI Innovation

Our analysis centers on the role of physical AI innovation — the absorption of AI into non-IT innovation — in firm-level production. At the firm level, we define *PhysicalAI* as the percentage

[18] Since the dependent variable is the natural log of ξ, and I_{AIcite} is an indicator variable, the coefficient 0.0039 in column 4 represents $(e^{0.0039}-1)*100=$ 0.4% higher market value.

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.[19] 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.[20]

Table 4 summarizes *PhysicalAI* and $I_{PhysicalAI}$, as well as all of the production input and output variables and controls. 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 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 – we find that 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.

III. Methodology

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 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:

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

---

[19] 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.

[20] 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.

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, our primary outcome is gross profit — revenues minus cost of goods sold (COGS). We then decompose its drivers by examining output (revenues plus the change in inventories) and cost efficiency (COGS divided by output).

III.a Jordà (2005) Local Projection

We follow Kogan et al. (2017), Kelly, Papanikolaou, Seru, and Taddy (2021) and Baslandze et al. (2025) in estimating the effect of *PhysicalAI* on production using the Jordà (2005) Local Projection (LP) methodology. We estimate firm i's response to *PhysicalAI* at horizon τ according to:

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

The dependent variable is the τ-year log change in Y, where Y is alternatively the number of employees, the capital stock, revenue TFP, gross profits, output, and cost efficiency, and $\tau = 1, \dots, 5$ years. We condition on the pre-period level of outcome by including the lagged dependent variable, $Y_{i,t-1}$. We control for firm size and risk dynamics by including the concurrent log value of capital stock and employment (when that variable is not the dependent variable) and log idiosyncratic volatility in vector X_{it}. We also include the log of one plus the number of patents filed in the year as a measure of innovative capacity. Finally, we isolate the dynamic control group with the industry-year by including SIC2 times year fixed effects δ_{jt}, double clustering standard errors by firm and year.

III.b LP Instrumental Variables: Laboratory AI

One might argue that effects identified in the LP analysis reflect firms' accumulated knowledge capital and endogenous decision-making rather than physical AI innovation. For example, firms with deeper expertise in adjacent technologies may be both more likely to engage in *PhysicalAI* and be better positioned to translate it into production growth. We address these concerns by implementing the linear projection-instrumental variable (LP-IV) methodology of Ramey (2016).

We 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 is motivated by Jaffe, Trajtenberg, Henderson (1993) and Helper et al. (2025), who document that knowledge spillovers are highly localized geographically**,** with firms disproportionately citing patents produced nearby. Universities and federal laboratories generate frontier scientific knowledge that is plausibly exogenous to individual firms' production and employment decisions, particularly among the non-IT firms in our sample. A higher local stock of AI patents and AI-absorbing patents from these institutions increases a firm's exposure to AI knowledge through localized spillovers, providing a source of variation in *PhysicalAI* that is independent of firm-specific capabilities or strategies.

We calculate the local AI knowledge stock (*Laboratory AI*) for patent p filed in year t as the total number of AI patents or AI-citing patents granted to universities and federal laboratories located in the same commuting zone as the inventor of patent p over the five years from *t-5* to *t-1*. If a patent has multiple inventors, we use the highest value of local AI stock among its inventors. We then collapse the patent level measure, local AI stock, to the firm-level by taking the mean value of this across all patents in a firm-year.

Our exclusion restriction requires that, conditional on region and industry controls, the historical accumulation of AI patents in nearby academic and government institutions is unrelated to contemporaneous firm-specific shocks to production inputs or outcomes except through the firm's AI knowledge absorption.[21] A potential threat to the exclusion restriction is that local AI knowledge stocks may be correlated with time-varying commuting zone characteristics - local economic dynamism, agglomeration effects, or industry composition shifts - that affect firm production independently of AI knowledge spillovers. For just over 40% of patents, however, the inventor is located in a different commuting zone than the assignee firm, generating variation in Laboratory AI across firms with headquarters in the same commuting zone. We can include CZ times year fixed effects, $\psi_{CZ,t}$ into the equation (2) specification, retaining firm-level controls and SIC2 times year fixed effects, $\delta_{j,t}$.

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

The inclusion of $\psi_{CZ,t}$ forces identification to the following logic. Consider two farm machinery firms, A and B, headquartered in the same commuting zone, say in Kansas City, with the same level of patenting. Firm A has a patenting employee located in the San Francisco Bay Area, proximate to AI expertise, while Firm B does not. To the econometrician, Firm A has a higher Laboratory AI value, with the scale depending on how much of the firm's innovation is invented in the Bay versus Kansas City. Now consider two precision agriculture firms, C and D, both headquartered in the Bay Area with equivalent patenting. Firm C focuses on farm equipment integration and has an inventor in Kansas City working with tractor partners whose patents do not reference AI; Firm D does not. Firm C accordingly has lower Laboratory AI exposure than Firm D. In this setup, a violation of the exclusion restriction would require that the inventor locations in these examples correlate with general R&D spillovers or research talent networks that independently affect firm outcomes — after controlling for patenting activity, firm size, commuting zone × year fixed effects, and industry × year fixed effects. While violations are possible, we consider them unlikely given the stringency of this comparison. Moreover, the general knowledge spillovers most likely to accompany inventor location are precisely those tied to why the inventor is there — which is the firm's physical AI work itself.

Nevertheless, we take two further steps toward confidence that the exclusion restriction holds. First, we run a placebo test to look for violations (presented in Appendix Table A6). If our results are driven by general rather than AI-specific knowledge spillovers, Laboratory AI should directly predict firm outcomes even among innovative firms that do not engage in physical AI innovation. We find no such effect: i.e., the relationship between firm outcomes and Laboratory AI is statistically insignificant across all horizons, consistent with the channel being specifically AI knowledge absorption.

Second, we consider the possibility that inventor location systematically differs from firm headquarters in ways that could compromise the exclusion restriction. If inventors are stationed at the firm's primary operating locations rather than at headquarters, the Laboratory AI instrument may reflect local economic conditions at those sites—conditions that our headquarter-based commuting zone × year fixed effects would not absorb. To address this, we follow the text-analysis strategy of Garcia and Norli (2012) to identify each firm's primary operating state as the state most frequently referenced across 10-K Items 1, 2, 6, and 7. We then document the robustness of our results to using primary operating state × year fixed effects

Figure 3 maps Laboratory AI intensity across U.S. commuting zones. To illustrate the sources of variation, we highlight three commuting zones, circled in red on the map. The first (the one in the middle of the map) is the Iowa City / Davenport commuting zone. It has a moderate level of AI, coming primarily from the University of Iowa, although there is a U.S. Geological Survey (USGS) – Central Midwest Water Science Center locally. Hence the university is driving the measure. The second, the San Francisco Bay Area, is among the most AI-intensive zones in the country. AI knowledge here emerges from UC Berkeley, UCSF, UC Davis, and other universities. These universities have large AI patenting. Yet, this commuting zone is exceptional also because it houses two major Department of Energy (DOE) laboratories – the Lawrence Livermore National Laboratory and the Lawrence Berkeley National Laboratory. The third, the Washington D.C. metro area (the "DMV" — District, Maryland, Virginia), tells a different story. While the region hosts notable research universities including Georgetown, the University of Maryland, George Washington, George Mason, and the Naval Academy, Laboratory AI here is also saturated by federal research institutions: NIH, NIST, NASA, and a dense concentration of DOE and Department of Defense (DOD) facilities.

IV. Results

The results tables all flow with a similar cadence as follows. In each table, panel A presents the linear projection (LP) results, and panel B reports the linear projection-instrumental variables (LP-IV) results. Within the panels, columns correspond to successive projection horizons, tracing the cumulative response of each outcome variable over time.

IV.a Production Factors – Employment, Capital, and TFP

We begin with the empirical findings that study the factors of production and their productivity as dependent variables. Table 5, Panel A presents LP estimates for employment. Physical AI innovation has a positive and statistically significant effect on employment across all reported horizons. Because these are cumulating linear projections, we might expect the coefficients to be increasing over time (because the years cumulate) but at a decreasing rate (because the effect diminishes in impact far into the horizon). This is what we find, a concave projection, significant through our horizon.

Panel B reports the equivalent to Panel A, but with *PhysicalAI* instrumented with *Laboratory AI* and including commuting zone times time fixed effects. We first consider the first-stage for relevance. The reported Kleibergen-Paap F-statistics are higher than 95 across all horizons in Table 5. These first stage results as well as the reported, but not discussed, F-statistics in subsequent tables are well above conventional weak-instrument thresholds, confirming the strong relevance of Laboratory AI as an instrument for *PhysicalAI*.

In the second stage, the coefficient on instrumented *PhysicalAI* is positive and statistically significant in explaining changes in log employment. The *PhysicalAI* effect appears across most of the 5-year horizon, with some weakening statistical confidence over the cumulating time. The effect has the same increasing but concave effect. Note that the magnitude cumulating to year 5 is about double that of the first year effect. As indicated in the methodology, we replicate these results using firms' primary operating state times year fixed effects (plus the fixed effect for the headquarter commuting zone) in lieu of the headquarters commuting zone times year fixed effects. Appendix Table A7 reports these results for employment in Panel A and includes a panel for each of the dependent variables studied in the remainder of the paper. The results are unchanged, which alleviates concerns that our instrument is simply capturing economic conditions at the firms' operating locations.

Before converting these estimates to economic magnitude inferences, we address the magnitude gap between panels A and B: the LP-IV coefficient is approximately six times larger than the LP coefficient. While attenuation bias in the LP could partly account for this gap, we favor a more conservative interpretation. We turn to Jiang (2017), who identifies three sources of what she calls "implausibly large IV estimates." The first is specification search - iterating on the instrument until a large magnitude emerges, the IV analog of p-hacking. Our incentive runs in the opposite direction: a smaller IV magnitude would be more convenient, not a larger one. The second cause according to Jiang (2017) is weak instruments, which is not our setting given first-stage F-statistics.

The third possibility for larger IV estimates - and the one Jiang finds most prevalent in top-tier finance publications - is that IV estimates recover a local average treatment effect (LATE) rather than a population average treatment effect (ATE). Drawing on Card (2001) and Angrist and Pischke (2009), she illustrates this with the college proximity instrument: the students most affected by education are precisely those induced to attend college by proximity, so the LATE can exceed the ATE. Our setting is analogous. The firms most affected by physical AI are likely those that have actively sought out AI expertise - for instance, by locating inventors in high-Laboratory-

AI commuting zones to intentionally absorb local AI knowledge. Jiang (2017) shows that under heterogeneous treatment effects, IV estimates can diverge further from the population average treatment effect than OLS, even when the exclusion restriction holds. The instrument remains valid, but the magnitude is more accurately interpreted as a LATE than an ATE.

To be conservative, we interpret magnitudes as a range spanning the LP and LP-IV estimates. Following Jiang (2017), this range reflects heterogeneity in treatment effects across locations rather than a robustness check for endogeneity.

Interpreting the coefficients in economic magnitudes, we find that a standard deviation (0.067 from Table 4) higher *PhysicalAI* results in a 0.28% (= 0.041 x 0.067) higher employment growth one year forward in the LP estimation (panel A) and 2.57% higher employment growth in the LP-IV (panel B). Perhaps the more interesting result is that the magnitude of the employment effect increases through horizon year 5. Hence, applying the panel A and B coefficients from Table 5 at t+5 and converting to average annual growth, we find that a standard deviation more physical AI innovation results in a 0.16% to 1.02% growth in employment annually, with the range reflecting location-specific treatment effects. 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. [22]

Table 6 reports an analogous positive and significant relationship between physical AI innovation and capital stock. The fact that both labor and capital respond to physical AI is reassuring in being consistent with factors of production being complementary. The capital effects, however, do not persist as long as those for employment, which may reflect that capital investment is front-loaded while employment continues to build out as AI-absorbed innovations are deployed. Both the LP and LP-IV estimates remain significant through the three-year horizon, after which the LP estimate loses significance; we therefore focus on t+3. At this horizon, a one-standard-deviation increase in physical AI innovation implies annualized capital stock growth of 0.17% to 2.13%, based on the Panel A and Panel B coefficients from Table 6.

[22] In Appendix Table A5, we divide the LP result into process (35% of patents) and product (65% of patents) *PhysicalAI*. Even though Table A4 already indicated that *PhysicalAI* is not simply a proxy for previously studied process innovation, we are interested in the employment effect in both product and process *PhysicalAI*. We find that product *PhysicalAI* and process *PhysicalAI* both project increased employment growth over the full horizon, consistent with Baslandze et al. (2025).

Table 7 reports the TFP results. In the LP specification, *PhysicalAI* shows no positive effect on TFP. The coefficient is marginally significant and negative at the three-year horizon, suggesting that cross-sectional variation in physical AI does not map to productivity gains on average. The LP-IV tells a different story: positive and significant effects emerge at the two-, three-, and five-year horizons, though the pattern is non-monotonic with an insignificant coefficient at t+4. The five-year estimate implies annualized TFP growth of 1.46%. As with employment and capital, this magnitude reflects a LATE rather than an economy-wide average effect.

Our results indicate that physical AI innovation is valued by the market and drives expansion in both factors of production — employment and capital. Productivity gains are present in the LP-IV but not robust across methods. This 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 or through realized variable production cost efficiencies – all potentially making use of the additional capital and labor we document. Next, we explore these mechanisms through income statement measures of profits, revenues, and costs.

IV.b Production Outcomes – Profits, Revenues, and Cost Efficiency

Tables 8, 9 and 10 present LP and LP-IV estimates where the dependent variables are production outcomes rather than inputs: log of gross profits, log of output (revenues plus changes in inventories), and COGS/Output as cost efficiency. The specifications mirror those of Tables 5-7, with covariates including the lagged log dependent variable, log idiosyncratic volatility, log of one plus the lagged patent count, and SIC2 times year fixed effects. If physical AI innovation drives growth in labor and capital, as our earlier results suggest, we would expect these input expansions to manifest eventually in profits. We can then investigate the channels: does physical AI innovation enhance revenue, or does it improve production efficiency by reducing variable input costs relative to output? A physical AI effect operating through cost efficiency - lower COGS/output - would represent a channel distinct from revenue TFP, which captures productivity as a residual from the production function rather than through cost-side improvements.

We find strong evidence that physical AI innovation impacts forward growth in profits, continuing through the 5-year projection horizon. The coefficients on *PhysicalAI* in both panels A and B of Table 8 exhibit consistently strong statistical significance. At a five-year horizon, a one standard deviation larger *PhysicalAI* implies that profits grow at annualized rates in the range of 0.22% to 1.80%, comparable to the annualized economic magnitude of the *PhysicalAI* on employment (ranging from 0.16% to 1.02%) and capital stock (0.17% to 2.13%).

Tables 9 and 10 provide evidence as to the origins of the physical AI profit effect. Profit can be thought of as revenues minus costs. In Table 9, we study revenues using the standard output measure of sales plus change in inventory. The LP regressions show a robust positive relation between physical AI innovation and output growth across all horizons. In the LP-IV specification, however, the output results are weaker and less uniform in statistical significance. We do find evidence of positive revenue growth, but at the two- and five-year cumulative horizons. The economic magnitudes are consistent with those from our employment and capital results. At a five-year horizon, a one-standard-deviation increase in *PhysicalAI* implies higher output growth at annualized rates of 0.17% and 1.22% from Panels A and B, respectively.

The cost efficiency results in Table 10 are unambiguous. We find no evidence that physical AI innovation affects COGS/output in any horizon in either panel. The patterns across Tables 8–10 suggest that the profit gains from physical AI innovation flow through revenue expansion rather than input cost reduction. 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.c Smart Automation Results

The literature on AI and labor markets has devoted substantial attention to automation.[23] 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.[24]

Table 11 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 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 Table 10: 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.[25] Panel A of Table 11 shows that both *SmartAutomation* and *NonAutoPhysicalAI* positively predict profit growth, with *NonAutoPhysicalAI* somewhat more consistently significant across horizons. The similarities end there. Panel C 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

---

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

[24] 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.

[25] Because we lack an instrument for the interacted specification of PhysicalAI with automation status, we limit this analysis to the LP framework.

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 12, 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 11 and 12 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 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. Using the patent value framework of Kogan et al. (2017), we show that physical AI patents are more valuable than comparable patents, consistent with the idea that combining AI with domain-specific technologies generates economically meaningful innovations.

To understand the higher value of physical AI patents, we examine firm production inputs and outcomes. Across several empirical approaches—including local projection estimation and an instrumental-variable strategy based on localized AI knowledge spillovers—we find that firms engaging in physical AI innovation experience higher employment and capital growth. We also find some evidence, albeit less robust, that physical AI improves productivity. Looking at production outcomes, we find that physical AI innovation predicts higher growth in profits and outputs but not cost efficiency. 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.

References

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

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

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

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

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

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

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

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

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

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

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

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.

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.

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

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

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

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

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

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

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

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

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

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

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

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

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen. 2025. "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. 2025. "Generative AI at Work*." *The Quarterly Journal of Economics* 140 (2): 889–942.

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

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

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

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

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

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

Hall, Bronwyn H., Adam 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.

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

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

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

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

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

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

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

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

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

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

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

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

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

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.

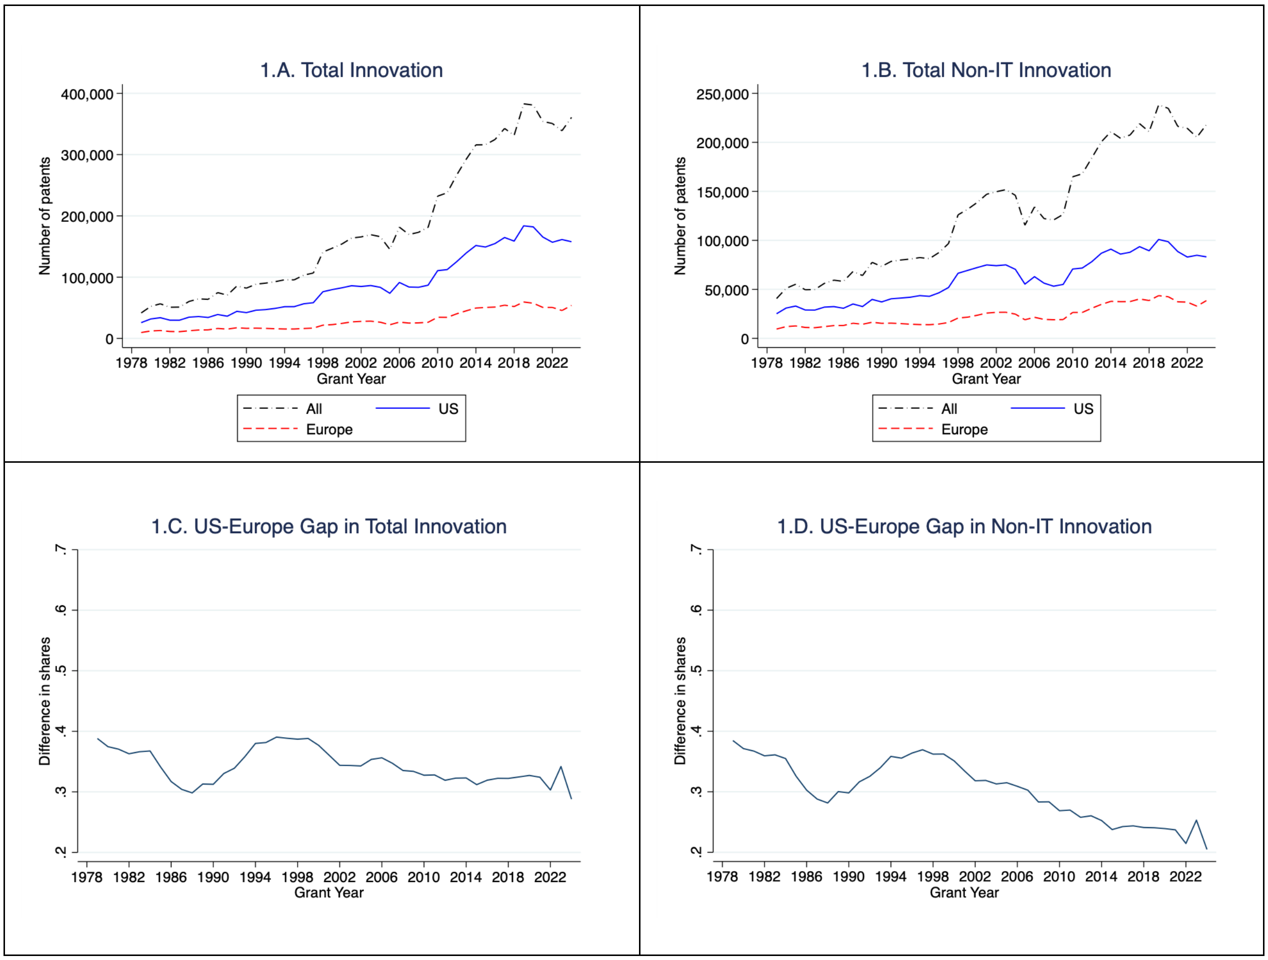

Figure 2: AI Innovation and Physical AI Innovation

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.

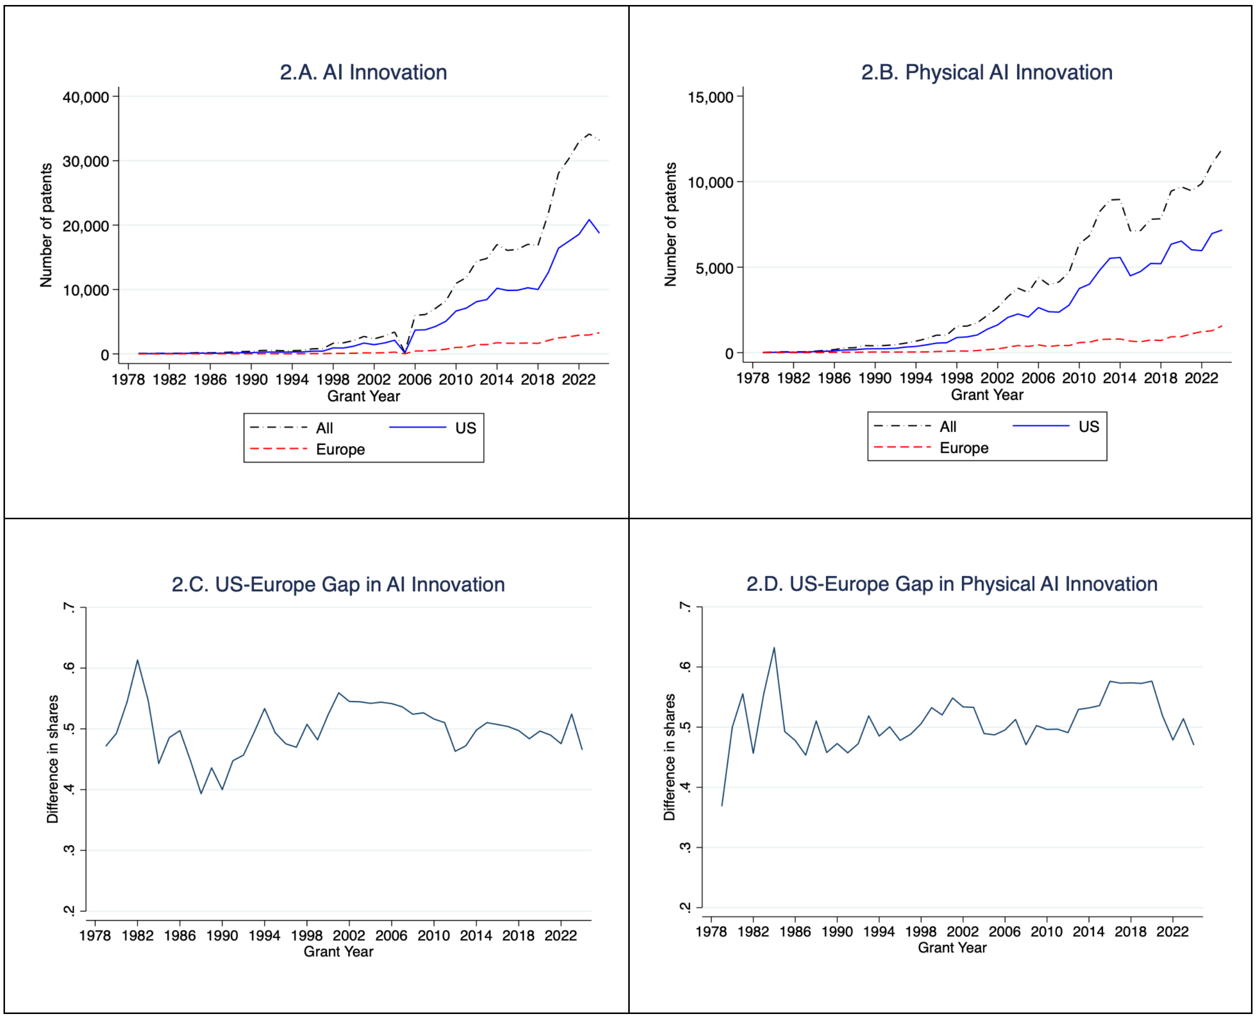

Figure 3: Heat Map of Laboratory AI Instrument

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.

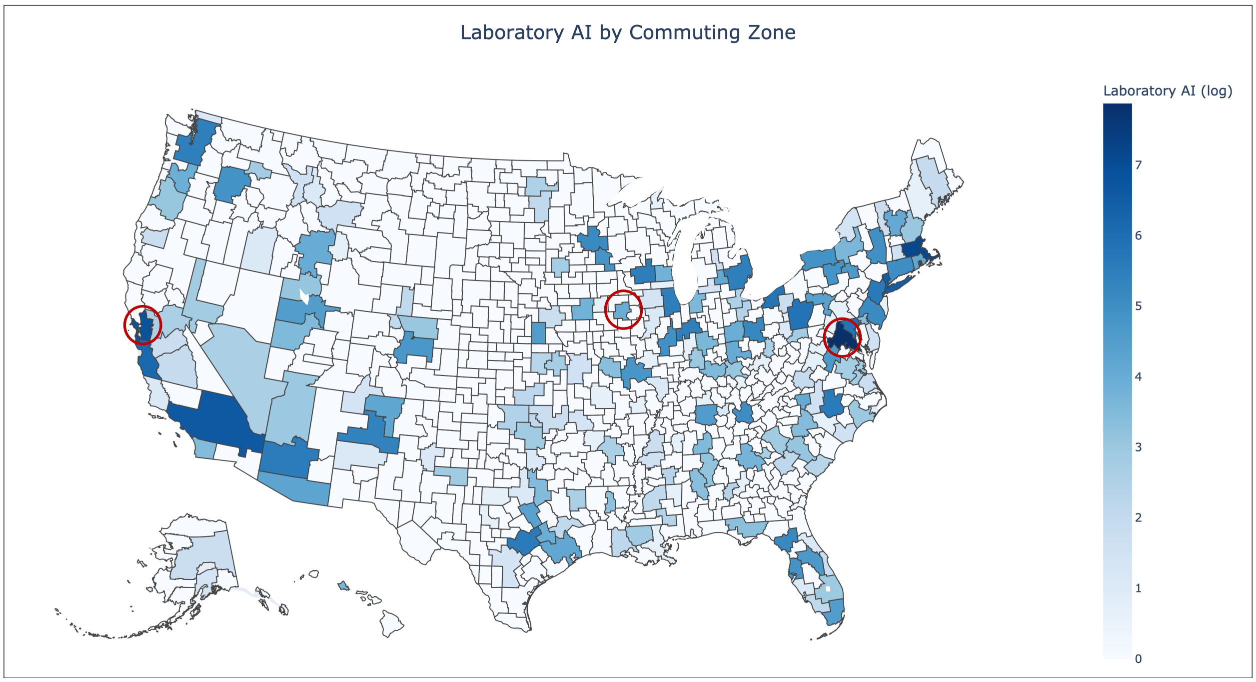

Table 1: Industry of Firms with Physical AI Innovation

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
Originality	2489656	0.394	0.292	0.446	0	0.973
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
Originality	1,722,139	0.405	0.292	0.469	0	0.973
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:

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

In columns 1 and 3, V is patent originality defined as 1 minus the HHI of the technology-class distribution of patent p's backward citations following Hall, Jaffe, and Trajtenberg (2001). In columns 2 and 4, V is $ln(\xi_p)$ 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. ***, **,* indicate significance at the 1%, 5%, and 10% respectively. All variable definitions are provided in Appendix A. The sample period is 1979 to 2024.

	IT Patents		Non-IT Patents	
	(1)	(2)	(3)	(4)
	Originality	$ln(\xi)$	Originality	$ln(\xi)$
I_{AIcite}	0.0128***	0.0014	0.0626***	0.0039**
	(6.65)	(1.32)	(19.00)	(2.15)
ln Market Cap	-0.0031*	0.8679***	-0.0022**	0.8900***
	(-1.70)	(44.70)	(-2.07)	(78.88)
ln Backward	0.1154***	0.0017***	0.1174***	0.0014***
	(72.49)	(3.18)	(104.62)	(3.31)
Observations	716,455	716,823	1,567,386	1,605,471
R-squared	0.3770	0.9625	0.3712	0.9638
Firm x Year FE	Yes	Yes	Yes	Yes
Patent Class Fe	Yes	Yes	Yes	Yes
Clustered SE	Firm x Yr	Firm x Yr	Firm x Yr	Firm x Yr

Table 4: Firm-level Summary Statistics

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

Variable	N	Mean	SD	Median	Min	Max
Physical AI Innovation	263,248	0.007	0.067	0	0	1
$I_{PhysicalAI}$	263,248	0.030	0.170	0	0	1
ln Employment	215,887	-0.167	2.340	-0.124	-6.908	6.153
ln Capital Stock	220,498	3.750	2.805	3.649	-6.999	11.211
Revenue TFP	145,869	0.002	0.419	0.008	-7.961	5.552
ln Profit	212,161	3.672	2.240	3.668	-8.066	10.292
ln Output	203,349	4.654	2.442	4.723	-7.918	10.777
COGS/Output	204,334	0.982	3.605	0.652	-0.242	99.735
ln (1+Patents)	237,124	0.425	1.018	0	0	8.092
ln(σ)	250,538	-3.615	0.625	-3.649	-5.307	-1.673
ln Laboratory AI	263,248	0.414	1.147	0	0	6.417

Table 5: Physical AI Innovation and Employment

This table presents the relation between physical AI innovation (denoted *PhysicalAI*) and employment. In Panel A, we present the coefficients β_τ for cumulative forward time periods $\tau \in \{1, \dots, 5\}$ from estimations of the following:

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

where Y is employment, and *Physical*AI is the fraction of the firm's patents that are AI-citing. 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+$\text{Patents}_{i,t-1}$). δ_{jt} are SIC2 times year fixed effects. In Panel B, we present estimates from the following instrumental variable estimation:

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

where $\widehat{Physical}$AI is the fitted value from a first stage estimation using the IV of *LaboratoryAI* -- the number of AI patents or AI-citing patents granted to universities and federal labs located in the firm's commuting zone. Also included in this specification, but not in the OLS are commuting zone times year fixed effects. In both panels, 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= Employment	Horizon				
	t+1	t+2	t+3	t+4	t+5
Panel A: Linear Projection					
Physical AI Innovation	0.041***	0.063***	0.089***	0.113***	0.117**
	(3.66)	(3.09)	(3.52)	(3.08)	(2.56)
ln (1+Patents) (lag)	0.006***	0.012***	0.018***	0.023***	0.027***
	(5.73)	(6.32)	(6.56)	(6.62)	(6.45)
ln Employment (lag)	-0.037***	-0.062***	-0.082***	-0.099***	-0.116***
	(-18.02)	(-20.09)	(-19.25)	(-17.69)	(-16.57)
Observations	164385	148238	134035	121597	110643
R-squared	0.057	0.073	0.083	0.091	0.100
Industry x Year FE & Controls	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year
Panel B: IV-Linear Projection					
First Stage					
Laboratory AI	0.010***	0.010***	0.010***	0.010***	0.010***
First stage F-statistic	123.30	114.67	107.97	100.96	95.70
Second Stage					
$\widehat{PhysicalAI}$	0.383***	0.393**	0.461	0.636*	0.762*
	(3.59)	(2.02)	(1.57)	(1.74)	(1.69)
ln (1+Patents) (lag)	0.004**	0.010***	0.016***	0.021***	0.025***
	(2.68)	(4.28)	(4.46)	(4.45)	(4.36)
ln Employment (lag)	-0.039***	-0.065***	-0.085***	-0.106***	-0.122***
	(-17.27)	(-19.43)	(-18.23)	(-16.88)	(-15.52)
Observations	140592	127019	114951	104297	94863
Industry x Year FE & Controls	Yes	Yes	Yes	Yes	Yes
Commuting Zone x Year FE	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year

Table 6: Physical AI Innovation and Capital Stock

This table presents the relation between physical AI innovation (denoted *PhysicalAI*) and capital stock. In Panel A, we present the coefficients β_τ for cumulative forward time periods $\tau \in \{1, \ldots, 5\}$ from estimations of the following:

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

where Y is capital stock, and *Physical*AI is the fraction of the firm's patents that are AI-citing. 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 employment, and ln(1+$Patents_{i,t-1}$). δ_{jt} are SIC2 times year fixed effects. In Panel B, we present estimates from the following instrumental variable estimation:

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

where $\widehat{Physical}$AI is the fitted value from a first stage estimation using the IV of *LaboratoryAI* -- the number of AI patents or AI-citing patents granted to universities and federal labs located in the firm's commuting zone. Also included in this specification, but not in the OLS, are commuting zone times year fixed effects. In both panels, 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= Capital Stock	Horizon				
	t+1	t+2	t+3	t+4	t+5
Panel A: Linear Projection					
Physical AI Innovation	0.044***	0.054*	0.076*	0.083	0.075
	(2.87)	(1.82)	(1.98)	(1.57)	(1.14)
ln (1+Patents) (lag)	0.010***	0.018***	0.024***	0.030***	0.037***
	(5.72)	(6.42)	(6.11)	(6.29)	(6.29)
ln Capital Stock (lag)	-0.061***	-0.099***	-0.128***	-0.150***	-0.171***
	(-13.48)	(-13.68)	(-14.04)	(-14.47)	(-15.55)
Observations	167201	150968	136597	123964	112783
R-squared	0.078	0.099	0.115	0.124	0.132
Industry x Year FE & Controls	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year
Panel B: IV-Linear Projection					
First Stage					
Laboratory AI	0.010***	0.010***	0.010***	0.010***	0.010***
First stage F-statistic	122.47	116.02	108.04	99.59	93.65
Second Stage					
$\widehat{PhysicalAI}$	0.740***	1.054***	0.952**	1.073*	0.900
	(4.64)	(3.35)	(2.15)	(1.72)	(1.33)
ln (1+Patents) (lag)	0.005*	0.011**	0.017***	0.022***	0.030***
	(1.99)	(2.69)	(3.28)	(3.24)	(3.68)
ln Capital Stock (lag)	-0.065***	-0.104***	-0.133***	-0.155***	-0.175***
	(-15.06)	(-15.10)	(-15.23)	(-15.47)	(-16.45)
Observations	142514	128851	116674	105894	96296
Industry x Year FE & Controls	Yes	Yes	Yes	Yes	Yes
Commuting Zone x Year FE	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year

Table 7: Physical AI Innovation and TFP

This table presents the relation between physical AI innovation (denoted *PhysicalAI*) and revenue total factor productivity (TFP). In Panel A, we present the coefficients β_τ for cumulative forward time periods $\tau \in \{1, \ldots, 5\}$ from estimations of the following:

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

where Y is TFP, and *Physical*AI is the fraction of the firm's patents that are AI-citing. 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 employment and capital stock, and ln(1+Patents$_{i,t-1}$). δ_{jt} are SIC2 times year fixed effects. In Panel B, we present estimates from the following instrumental variable estimation:

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

where $\widehat{Physical}$AI is the fitted value from a first stage estimation using the IV of *LaboratoryAI* -- the number of AI patents or AI-citing patents granted to universities and federal labs located in the firm's commuting zone. Also included in this specification, but not in the OLS, are commuting zone times year fixed effects. In both panels, 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= TFP	Horizon				
	t+1	t+2	t+3	t+4	t+5
Panel A: Linear Projection					
Physical AI Innovation	0.001	-0.021	-0.026*	-0.006	-0.009
	(0.11)	(-1.46)	(-1.91)	(-0.30)	(-0.31)
ln (1+Patents) (lag)	0.000	0.003*	0.006***	0.009***	0.011***
	(0.17)	(1.69)	(3.19)	(4.19)	(4.48)
ln TFP (lag)	-0.235***	-0.377***	-0.481***	-0.578***	-0.644***
	(-23.32)	(-32.22)	(-30.21)	(-38.14)	(-39.72)
Observations	114202	102183	92439	83891	76343
R-squared	0.075	0.121	0.162	0.207	0.235
Industry x Year FE & Controls	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year
Panel B: IV-Linear Projection					
First Stage					
Laboratory AI	0.010***	0.010***	0.010***	0.010***	0.010***
First stage F-statistic	98.52	88.09	84.07	76.57	71.30
Second Stage					
$\widehat{PhysicalAI}$	0.158	0.688**	0.636**	0.424	1.087**
	(0.86)	(2.54)	(2.03)	(1.22)	(2.57)
ln (1+Patents) (lag)	-0.001	-0.002	0.002	0.006**	0.003
	(-0.37)	(-0.64)	(0.76)	(2.27)	(0.98)
ln TFP (lag)	-0.238***	-0.374***	-0.479***	-0.579***	-0.645***
	(-24.48)	(-29.62)	(-28.78)	(-36.82)	(-39.17)
Observations	97533	87466	79198	71922	65409
Industry x Year FE & Controls	Yes	Yes	Yes	Yes	Yes
Commuting Zone x Year FE	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year

Table 8: Physical AI Innovation and Profits

This table presents the relation between physical AI innovation (denoted *PhysicalAI*) and profits. In Panel A, we present the coefficients β_τ for cumulative forward time periods $\tau \in \{1,\ldots,5\}$ from estimations of the following:

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

where Y is gross profits (sales minus COGS), and *Physical*AI is the fraction of the firm's patents that are AI-citing. 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 employment and capital stock, and ln(1+$Patents_{i,t-1}$). δ_{jt} are SIC2 times year fixed effects. In Panel B, we present estimates from the following instrumental variable estimation:

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

where $\widehat{Physical}$AI is the fitted value from a first stage estimation using the IV of *LaboratoryAI* -- the number of AI patents or AI-citing patents granted to universities and federal labs located in the firm's commuting zone. Also included in this specification, but not in the OLS, are commuting zone times year fixed effects. In both panels, 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=Profit	Horizon				
	t+1	t+2	t+3	t+4	t+5
Panel A: Linear Projection					
Physical AI Innovation	0.060***	0.090***	0.122***	0.153***	0.161***
	(3.52)	(3.63)	(3.45)	(3.76)	(2.97)
ln (1+Patents) (lag)	0.010***	0.018***	0.024***	0.030***	0.036***
	(7.93)	(7.82)	(7.66)	(7.82)	(7.95)
ln Profit (lag)	-0.069***	-0.096***	-0.115***	-0.129***	-0.144***
	(-19.06)	(-20.59)	(-17.21)	(-15.67)	(-15.59)
Observations	151001	136828	124365	113339	103522
R-squared	0.090	0.109	0.116	0.122	0.132
Industry x Year FE & Controls	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year
Panel B: IV-Linear Projection					
First Stage					
Laboratory AI	0.013***	0.013***	0.013***	0.013***	0.013***
First stage F-statistic	115.78	109.55	106.44	98.51	89.71
Second Stage					
$\widehat{PhysicalAI}$	0.516***	0.751***	0.774***	1.002***	1.343***
	(5.17)	(3.59)	(2.90)	(3.18)	(3.52)
ln (1+Patents) (lag)	0.007***	0.013***	0.021***	0.025***	0.029***
	(3.84)	(4.18)	(5.15)	(4.77)	(4.67)
ln Profit (lag)	-0.072***	-0.100***	-0.122***	-0.137***	-0.155***
	(-18.63)	(-19.85)	(-17.34)	(-15.41)	(-16.17)
Observations	128675	116743	106165	96759	88301
Industry x Year FE & Controls	Yes	Yes	Yes	Yes	Yes
Commuting Zone x Year FE	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year

Table 9: Physical AI Innovation and Output

This table presents the relation between physical AI innovation (denoted *PhysicalAI*) and output. In Panel A, we present the coefficients β_τ for cumulative forward time periods $\tau \in \{1, \ldots, 5\}$ from estimations of the following:

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

where Y is output (sales plus the change in inventory deflated by the CPI), and *Physical*AI is the fraction of the firm's patents that are AI-citing. 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 employment and capital stock, and ln(1+$Patents_{i,t-1}$). δ_{jt} are SIC2 times year fixed effects. In Panel B, we present estimates from the following instrumental variable estimation:

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

where $\widehat{Physical}$AI is the fitted value from a first stage estimation using the IV of *LaboratoryAI* -- the number of AI patents or AI-citing patents granted to universities and federal labs located in the firm's commuting zone. Also included in this specification, but not in the OLS, are commuting zone times year fixed effects. In both panels, 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= Output	Horizon				
	t+1	t+2	t+3	t+4	t+5
Panel A: Linear Projection					
Physical AI Innovation	0.037**	0.057*	0.099**	0.138***	0.128**
	(2.41)	(2.01)	(2.52)	(2.84)	(2.05)
ln (1+Patents) (lag)	0.005***	0.012***	0.018***	0.026***	0.033***
	(3.98)	(5.07)	(5.95)	(6.55)	(7.24)
ln Output (lag)	-0.080***	-0.124***	-0.147***	-0.173***	-0.189***
	(-23.05)	(-21.03)	(-21.20)	(-18.76)	(-16.85)
Observations	146492	132119	120057	109327	99755
R-squared	0.075	0.101	0.106	0.113	0.121
Industry x Year FE & Controls	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year
Panel B: IV-Linear Projection					
First Stage					
Laboratory AI	0.011***	0.011***	0.011***	0.011***	0.011***
First stage F-statistic	115.30	107.74	100.42	94.56	86.12
Second Stage					
$\widehat{PhysicalAI}$	0.288	0.479*	0.398	0.461	0.907*
	(1.31)	(1.86)	(1.10)	(1.01)	(1.72)
ln (1+Patents) (lag)	0.003	0.008***	0.014***	0.021***	0.026***
	(1.62)	(2.74)	(3.62)	(4.37)	(4.32)
ln Output (lag)	-0.083***	-0.127***	-0.154***	-0.179***	-0.198***
	(-21.79)	(-20.43)	(-19.90)	(-17.79)	(-15.67)
Observations	125126	112976	102722	93541	85260
Industry x Year FE & Controls	Yes	Yes	Yes	Yes	Yes
Commuting Zone x Year FE	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year

Table 10: Physical AI Innovation and Cost Efficiency

This table presents the relation between physical AI innovation (denoted *PhysicalAI*) and cost efficiency. In Panel A, we present the coefficients β_τ for cumulative forward time periods $\tau \in \{1, \dots, 5\}$ from estimations of the following:

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

where Y is COGS/Output, and *Physical*AI is the fraction of the firm's patents that are AI-citing. 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 employment and capital stock, and ln(1+$Patents_{i,t-1}$). δ_{jt} are SIC2 times year fixed effects. In Panel B, we present estimates from the following instrumental variable estimation:

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

where $\widehat{Physical}$AI is the fitted value from a first stage estimation using the IV of *LaboratoryAI* -- the number of AI patents or AI-citing patents granted to universities and federal labs located in the firm's commuting zone. Also included in this specification, but not in the OLS, are commuting zone times year fixed effects. In both panels, 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: Linear Projection					
Physical AI Innovation	0.047	-0.045	-0.003	-0.197	-0.105
	(0.45)	(-0.39)	(-0.02)	(-1.25)	(-0.53)
ln (1+Patents) (lag)	0.016**	0.020**	0.021**	0.026**	0.021**
	(2.24)	(2.22)	(2.40)	(2.04)	(2.02)
ln COGS/Output (lag)	-0.083***	-0.162***	-0.174***	-0.280***	-0.271***
	(-3.14)	(-4.91)	(-4.36)	(-6.09)	(-4.33)
Observations	147888	133328	121099	110217	100535
R-squared	0.019	0.031	0.028	0.045	0.039
Industry x Year FE & Controls	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year
Panel B: IV-Linear Projection					
First Stage					
Laboratory AI	0.011***	0.011***	0.011***	0.011***	0.011***
First stage F-statistic	114.70	107.38	99.56	93.16	84.69
Second Stage					
$\widehat{PhysicalAI}$	-0.148	1.132	-0.750	0.756	1.168
	(-0.08)	(0.65)	(-0.31)	(0.34)	(0.43)
ln (1+Patents) (lag)	0.013	0.006	0.021	0.019	0.015
	(0.77)	(0.48)	(0.99)	(1.13)	(0.81)
ln COGS/Output (lag)	-0.089***	-0.155***	-0.193***	-0.292***	-0.282***
	(-5.10)	(-5.95)	(-3.84)	(-6.47)	(-5.34)
Observations	126303	113979	103580	94273	85902
Industry x Year FE & Controls	Yes	Yes	Yes	Yes	Yes
Commuting Zone x Year FE	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year

Table 11: Smart Automation

This table presents the relation between *Smart Automation* and firm-level outcomes. We present the coefficients β_τ for cumulative forward time periods $\tau \in \{1, \dots, 5\}$ 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: Profits				
Smart Automation (Automation Physical AI)	0.042*	0.038	0.040	0.111**	0.084
	(2.01)	(1.17)	(0.93)	(2.06)	(1.19)
Non-Automation Physical AI	0.052**	0.092***	0.138***	0.121**	0.150**
	(2.22)	(2.81)	(2.77)	(2.29)	(2.26)
Automation	0.026***	0.047***	0.061***	0.073***	0.085***
	(3.11)	(3.32)	(3.17)	(3.27)	(3.36)
Observations	151001	136828	124365	113339	103522
R-squared	0.090	0.109	0.117	0.122	0.132
	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 12: 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 11, 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: Product and Process Innovation as Physical AI Innovation

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

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

where Y are firm outcomes represented in the Panel names; $ProdPhysicalAI$ is the fraction of the firm's patents that are AI-citing product patents; $ProcessPhysicalAI$ is the fraction of the firm's patents that are AI-citing process patents. *Process* is the fraction of a firm's patents that are process patents. Process patents are patents with more than 50% process claims. The rest are classified as product patents. The classification of claims is described in Appendix C. X is a vector of firm-time level control variables including log value of capital stock, the log number of employees, log idiosyncratic volatility, $\ln(1+Patents_{i,t-1})$, and $\ln(Y_{i,t-1})$; δ_{jt} are SIC2 times year fixed effects. Standard errors are double-clustered by firm and year. ***, **,* indicate significance at the 1%, 5%, and 10% respectively. All variable definitions are provided in Appendix A. The sample period is 1979 to 2024.

	t+1	t+2	t+3	t+4	t+5
	Dependent Variable: ln Employment				
Product Physical AI	0.021*	0.045**	0.051**	0.066**	0.081**
	(1.83)	(2.69)	(2.14)	(2.20)	(2.22)
Process Physical AI	0.033**	0.067**	0.122***	0.120**	0.144**
	(2.32)	(2.43)	(3.12)	(2.33)	(2.44)
Process	0.011**	0.006	0.003	0.000	-0.008
	(2.66)	(0.80)	(0.25)	(0.04)	(-0.55)
Observations	164385	148238	134035	121597	110643
R-squared	0.057	0.073	0.083	0.092	0.101
Industry x Year FE	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year

Table A6: 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 A7: Physical AI Innovation and Firm Outcomes Under Alternate Fixed Effects

We present estimates from the following instrumental variable estimation:

$$\ln(Y_{i,t+\tau}) - \ln(Y_{i,t}) = \alpha_0 + \beta_\tau \widehat{PhysicalAI}_{i,t} + \gamma_1 X'_{i,t} + \delta_{j,t} + \vartheta_{s,t} + \phi_{CZ} + \epsilon_{i,t}$$

where Y are firm outcomes listed in each panel. $\widehat{Physical}$AI is the fitted value from a first stage regression of the firm's AI-citing patents on the instrument *LaboratoryAI* -- the number of AI patents or AI-citing patents granted to universities and federal labs located in the firm's commuting zone. 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, ϕ_{CZ} are commuting zone fixed effects based on the firm's headquarter location, and $\vartheta_{s,t}$ are primary state times year fixed effects. Primary state is the state most frequently mentioned in Items 1 (Business), Item 2 (Properties), Item 6 (Selected Financial Data), and Item 7 (Management Discussion and Analysis) of the firm's 10-K. 10-K filings are obtained from WRDS SEC Analytics Suite, which are available from 1993 onwards only. 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 1993 to 2024.

IV-Linear Projection

	Horizon				
	t+1	t+2	t+3	t+4	t+5
Panel A: Employment					
First Stage					
Laboratory AI	0.013***	0.013***	0.013***	0.013***	0.013***
First stage F-statistic	148.69	142.39	157.71	151.71	148.40
Second Stage					
$\widehat{PhysicalAI}$	0.470***	0.479***	0.459**	0.580**	0.594*
	(4.98)	(2.93)	(2.21)	(2.21)	(1.75)
ln (1+Patents) (lag)	0.002	0.008***	0.015***	0.018***	0.023***
	(1.58)	(3.60)	(4.57)	(4.32)	(4.29)
ln Employment (lag)	-0.040***	-0.065***	-0.085***	-0.104***	-0.119***
	(-13.89)	(-15.96)	(-15.43)	(-14.70)	(-13.83)
Observations	106313	96325	87239	79070	71767
Panel B: Capital					
First Stage					
Laboratory AI	0.013***	0.013***	0.013***	0.013***	0.013***
First stage F-statistic	149.95	145.58	157.37	151.83	147.37
Second Stage					
$\widehat{PhysicalAI}$	0.743***	1.090***	0.941***	1.001**	0.906*
	(4.76)	(4.09)	(2.77)	(2.23)	(1.79)
ln (1+Patents) (lag)	0.004	0.009**	0.017***	0.021***	0.027***
	(1.64)	(2.49)	(3.42)	(3.41)	(3.58)
ln Capital Stock (lag)	-0.064***	-0.106***	-0.137***	-0.162***	-0.182***
	(-12.03)	(-12.96)	(-12.97)	(-13.03)	(-13.31)
Observations	107388	97329	88198	79978	72562

Table A7 Continued

Panel C: TFP					
First Stage					
Laboratory AI	0.013***	0.012***	0.012***	0.012***	0.012***
First stage F-statistic	128.88	116.71	118.48	120.16	111.73
Second Stage					
$\widehat{PhysicalAI}$	0.138	0.588**	0.595*	0.558	0.925***
	(0.90)	(2.68)	(1.95)	(1.67)	(2.75)
ln (1+Patents) (lag)	0.000	-0.002	0.001	0.005	0.004
	(0.03)	(-0.82)	(0.21)	(1.62)	(1.22)
ln TFP (lag)	-0.239***	-0.375***	-0.488***	-0.591***	-0.650***
	(-20.65)	(-24.55)	(-23.96)	(-31.20)	(-35.96)
Observations	73401	66105	59986	54464	49468
Panel D: Profit					
First Stage					
Laboratory AI	0.017***	0.017***	0.016***	0.016***	0.016***
First stage F-statistic	141.01	135.93	146.08	141.71	135.97
Second Stage					
$\widehat{PhysicalAI}$	0.533***	0.856***	1.005***	1.111***	1.292***
	(5.26)	(5.46)	(4.59)	(4.41)	(4.62)
ln (1+Patents) (lag)	0.005***	0.011***	0.017***	0.022***	0.026***
	(3.01)	(3.77)	(4.34)	(4.10)	(4.16)
ln Profit	-0.071***	-0.103***	-0.122***	-0.133***	-0.149***
	(-16.59)	(-16.95)	(-14.87)	(-13.39)	(-13.40)
Observations	96438	87724	79851	72710	66223
Panel E: Output					
First Stage					
Laboratory AI	0.014***	0.014***	0.014***	0.014***	0.014***
First stage F-statistic	144.32	133.86	145.16	139.00	129.56
Second Stage					
$\widehat{PhysicalAI}$	0.335*	0.668**	0.546	0.877**	1.060**
	(1.73)	(2.54)	(1.62)	(2.36)	(2.55)
ln (1+Patents) (lag)	0.003	0.006*	0.012***	0.016***	0.023***
	(1.36)	(1.96)	(3.12)	(3.12)	(3.63)
ln Output (lag)	-0.083***	-0.130***	-0.157***	-0.184***	-0.200***
	(-19.86)	(-19.22)	(-17.74)	(-16.21)	(-14.08)
Observations	95462	86482	78697	71643	65176

Table A7 Continued

Panel F: COGS/Output					
First Stage					
Laboratory AI	0.014***	0.014***	0.014***	0.014***	0.014***
First stage F-statistic	136.78	131.59	143.19	137.56	128.46
Second Stage					
$\widehat{PhysicalAI}$	0.289	1.317	0.500	0.342	0.571
	(0.19)	(0.91)	(0.37)	(0.26)	(0.31)
ln (1+Patents) (lag)	0.015	0.009	0.020	0.034*	0.029
	(0.94)	(0.67)	(1.31)	(1.97)	(1.64)
ln COGS/Output (lag)	-0.093***	-0.147***	-0.185***	-0.295***	-0.280***
	(-4.48)	(-4.64)	(-3.24)	(-6.80)	(-5.45)
Observations	96421	87299	79392	72230	65686
Industry x Year FE & Controls	Yes	Yes	Yes	Yes	Yes
Commuting Zone x Year FE	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year

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 for the classification of IT patents and AI patents.

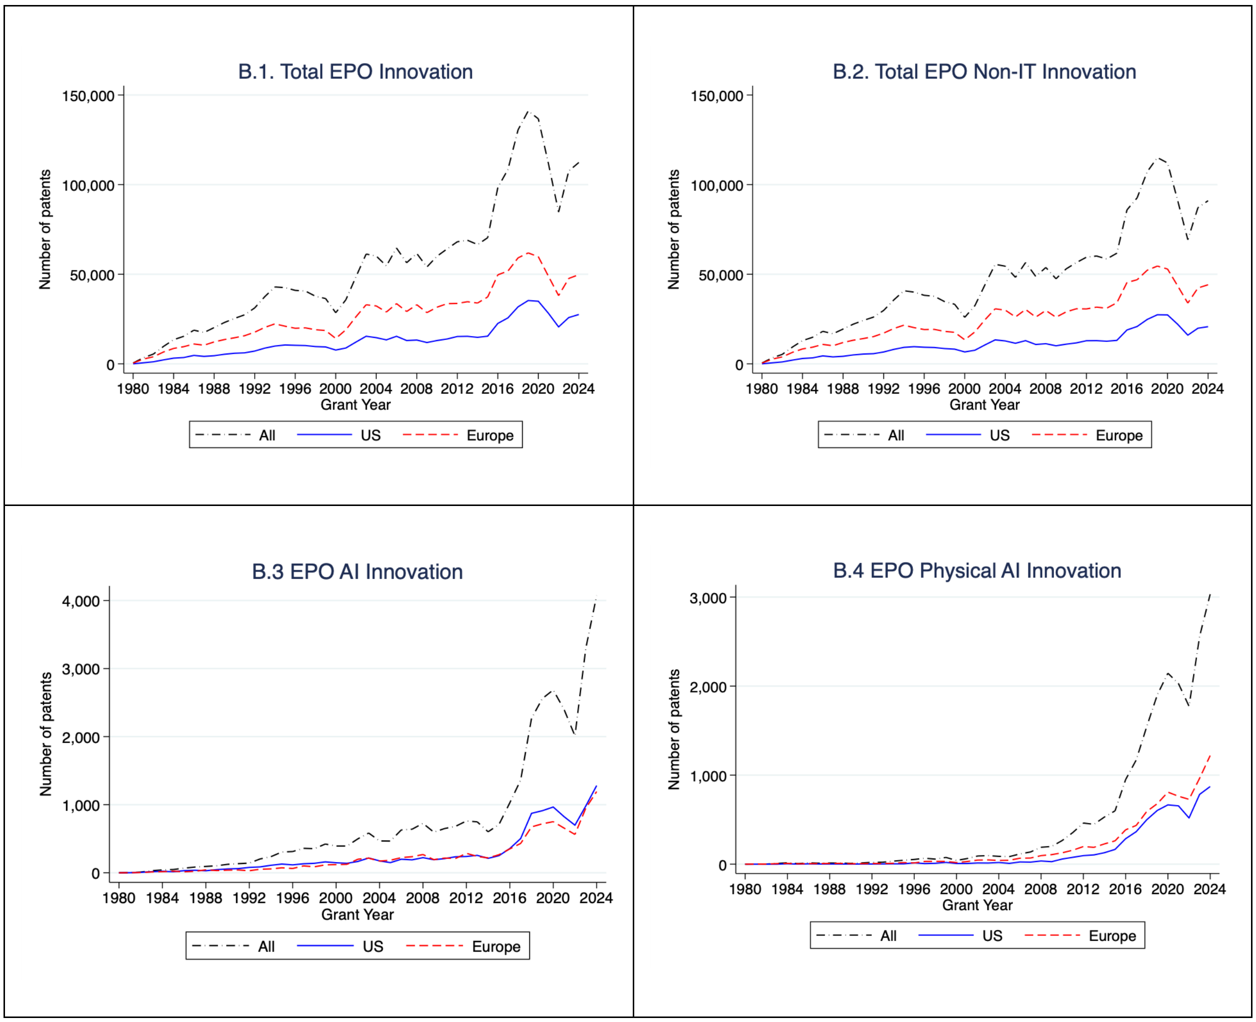

Appendix C: Automation Patents, Process Patents, and Cognitive/Abstract Occupations

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 characteristic.

Appendix D: Alternate Classification of AI Patents

In this section, we present robustness of our LP-IV results to a narrower definition of AI patents. Specifically, we exclude architectures necessary for AI applications such as parallel and distributed digital computing, including multi-processor configurations, interconnection networks, and array processing units. Table D1 below lists the IPC classes included in this narrower definition of AI. Table D2 presents results of our LP-IV estimation. We continue to find significant positive effects of physical AI innovation on employment, capital, TFP, profits, and output.

Table D1: 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

Table D2: Physical AI Innovation and Firm Outcomes Under Narrower AI Classification

We present estimates from the following instrumental variable estimation:

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

where Y are firm outcomes listed in each panel. $\widehat{Physical}$AI is the fitted value from a first stage regression of the firm's AI-citing patents on the instrument *LaboratoryAI* -- the number of AI patents or AI-citing patents granted to universities and federal labs located in the firm's commuting zone. X represents the same set of control variables used in main specifications shown in Tables 5 through 10. δ_{jt} 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.

IV-Linear Projection

	Horizon				
	t+1	t+2	t+3	t+4	t+5
Panel A: Employment					
First Stage					
Laboratory AI	0.013***	0.013***	0.013***	0.013***	0.013***
First stage F-statistic	143.54	137.62	152.49	151.94	150.45
Second Stage					
$\widehat{PhysicalAI}$	0.443***	0.452***	0.415*	0.501*	0.594*
	(4.74)	(2.80)	(1.93)	(1.90)	(1.71)
ln (1+Patents) (lag)	0.003**	0.010***	0.016***	0.022***	0.026***
	(2.45)	(4.39)	(5.11)	(5.06)	(4.99)
ln Employment (lag)	-0.039***	-0.065***	-0.085***	-0.106***	-0.122***
	(-17.30)	(-19.44)	(-18.22)	(-16.87)	(-15.51)
Observations	140592	127019	114951	104297	94863
Panel B: Capital					
First Stage					
Laboratory AI	0.013***	0.013***	0.013***	0.013***	0.013***
First stage F-statistic	145.75	140.84	152.47	151.20	148.79
Second Stage					
$\widehat{PhysicalAI}$	0.670***	0.910***	0.689**	0.659	0.560
	(4.52)	(3.65)	(2.09)	(1.48)	(1.13)
ln (1+Patents) (lag)	0.005**	0.012***	0.019***	0.025***	0.032***
	(2.51)	(3.55)	(4.39)	(4.30)	(4.54)
ln Capital Stock (lag)	-0.065***	-0.104***	-0.133***	-0.156***	-0.176***
	(-15.13)	(-15.20)	(-15.35)	(-15.61)	(-16.55)
Observations	142514	128851	116674	105894	96296

Table D2 Continued

Panel C: TFP

First Stage					
Laboratory AI	0.013***	0.012***	0.012***	0.012***	0.012***
First stage F-statistic	127.87	119.35	126.96	127.41	118.06
Second Stage					
$\widehat{PhysicalAI}$	0.142	0.563**	0.570**	0.579*	0.884***
	(1.03)	(2.58)	(2.05)	(1.97)	(2.85)
ln (1+Patents) (lag)	-0.001	-0.001	0.002	0.005**	0.005*
	(-0.39)	(-0.32)	(1.17)	(2.21)	(1.71)
ln TFP (lag)	-0.238***	-0.374***	-0.479***	-0.580***	-0.645***
	(-24.40)	(-29.56)	(-28.76)	(-36.93)	(-38.91)
Observations	97533	87466	79198	71922	65409

Panel D: Profit

First Stage					
Laboratory AI	0.016***	0.016***	0.016***	0.016***	0.016***
First stage F-statistic	141.75	135.99	145.42	142.16	136.83
Second Stage					
$\widehat{PhysicalAI}$	0.528***	0.781***	0.855***	1.006***	1.179***
	(5.17)	(4.72)	(4.02)	(3.99)	(4.02)
ln (1+Patents) (lag)	0.007***	0.013***	0.020***	0.025***	0.030***
	(4.08)	(4.61)	(5.40)	(5.03)	(5.12)
ln Profit	-0.073***	-0.101***	-0.122***	-0.137***	-0.154***
	(-18.45)	(-19.87)	(-17.43)	(-15.36)	(-16.01)
Observations	128675	116743	106165	96759	88301

Panel E: Output

First Stage					
Laboratory AI	0.014***	0.014***	0.014***	0.014***	0.014***
First stage F-statistic	141.93	132.61	141.78	137.70	129.59
Second Stage					
$\widehat{PhysicalAI}$	0.327*	0.630**	0.451	0.697*	0.911**
	(1.70)	(2.38)	(1.35)	(1.87)	(2.20)
ln (1+Patents) (lag)	0.003	0.007**	0.014***	0.019***	0.025***
	(1.68)	(2.54)	(3.90)	(4.13)	(4.51)
ln Output (lag)	-0.083***	-0.127***	-0.154***	-0.179***	-0.198***
	(-21.81)	(-20.31)	(-19.81)	(-17.97)	(-15.66)
Observations	125126	112976	102722	93541	85260

Table D2 Continued

Panel F: COGS/Output					
First Stage					
Laboratory AI	0.014***	0.014***	0.014***	0.014***	0.013***
First stage F-statistic	135.88	130.46	140.17	136.29	128.77
Second Stage					
$\widehat{PhysicalAI}$	-0.410	0.328	0.257	0.504	1.041
	(-0.27)	(0.26)	(0.19)	(0.42)	(0.69)
ln (1+Patents) (lag)	0.015	0.013	0.014	0.021*	0.016
	(1.14)	(1.39)	(1.05)	(1.70)	(1.40)
ln COGS/Output (lag)	-0.089***	-0.155***	-0.193***	-0.292***	-0.282***
	(-5.11)	(-5.97)	(-3.83)	(-6.45)	(-5.34)
Observations	126303	113979	103580	94273	85902
Industry x Year FE & Controls	Yes	Yes	Yes	Yes	Yes
Commuting Zone x Year FE	Yes	Yes	Yes	Yes	Yes
Clustered SE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year