Automation-Induced Innovation Shift *

Lin William Cong† Yao Lu‡ Hanqing Shi§ Wu Zhu¶

First draft: May 2024; current draft: June 2025.

Abstract

We study how exposure to automation affects the nature and level of corporate innovation, which informs how innovation begets innovation. We document that firms with high robot exposure alter their technological focus over time and shift innovative activities towards AI which automation naturally complements through data accumulation. The shift is more pronounced for firms with greater data generation or prior AI-related research experience. Because AI patents are more costly (e.g., in labor input), albeit more general and original, firms with high automation experience a significant rise in R&D expenditure but an initial drop in patent quantity, before benefiting. Our findings not only resolve the puzzle that globally firms invent less despite the greater research effort amidst rising automation, but also provides insights on the heterogeneous paths of innovation, all of which we rationalize in a parsimonious dynamic equilibrium model.

JEL Classification: G30, O36
Keywords: AI, Corporate Innovation, Language Models, Patents, Robot, Technology

*We are grateful to Daron Acemoglu, Tania Babina, Simba Xin Chang, Anastassia Fedyk, Pengfei Han, Zhiguo He, Gerard Hoberg, Po-Hsuan (Paul) Hsu, Leonid Kogan, Siguang Li, Josh Lerner, Evgeny Lyandres, Yueran Ma, Laura Veldkamp, Wenfeng Wu, Haotian Xiang, Liyan Yang, Yang You, Ben Miao Zhang, Hong Zhang, and Christina Zhu for their insightful comments and suggestions. We also thank participants at the 2025 AI and Climate Risk Forum at MIT — MRS International Risk Conference (Boston), 2025 Asian Financial Association Annual Meeting, 2025 ASSA Annual Meeting (Econometric Society), 2025 China International Conference in Finance (CICF), Five-Star Conference in Finance, as well as seminar participants at GSM Peking University, HKUST (GZ), and Tsinghua University for valuable feedback.

†Corresponding author, Cornell SC Johnson College of Business (Johnson), ABFER, & NBER, will.cong@cornell.edu. Mailing address: Samuel Curtis Johnson Graduate School of Management, Cornell University SC Johnson College of Business, Sage Hall, 114 Feeney Way, Ithaca, NY 14853-6201, USA. Phone: (607) 255-7859.

‡School of Economics and Management, Tsinghua University, luyao@sem.tsinghua.edu.cn

§School of Economics and Management, Tsinghua University, shq20@mails.tsinghua.edu.cn

¶School of Economics and Management, Tsinghua University, zhuwu@sem.tsinghua.edu.cn

Conflict-of-interest disclosure statement

Lin William Cong

I have nothing to disclose.

Yao Lu

Lu acknowledges financial support from the Humanities and Social Sciences Major Project of the Ministry of Education (No. 22JJD790047), the "Double-High" Program for Innovation Direction Construction in the Humanities and Social Sciences, and the "Impact Enhancement" Program of the School of Economics and Management at Tsinghua University (No. 2022051006).

Hanqing Shi

I have nothing to disclose.

Wu Zhu

Zhu acknowledges support from the NSFC (Grant No. 72273071), the Tsinghua University Initiative Scientific Research Program (No. 2022Z04W02016), and the School of Economics and Management Research Grant (No. 2022051002).

1 Introduction

The rapid diffusion of automation technologies—frequently dubbed the Fourth Industrial Revolution by industry leaders and the media—has fundamentally reshaped global production systems and sparked widespread public debate about its economic implications.[1] In particular, a growing body of research has documented the effects of automation on labor share, job displacement, wage dynamics, income inequality, and capital structure (e.g., Forbes, as well as Acemoglu and Restrepo 2020; Cheng et al. 2024). Yet, little is known about the implications of automation for corporate innovation—the very engine of long-run economic growth, as emphasized in Romer (1986), Aghion and Howitt (1992), and Mokyr (2011).

This article seeks to fill this gap by providing the first empirical evidence on how firm-level exposure to industrial robots and automation affects the direction, composition, and intensity of innovation over time. Using novel micro-level, text-based metrics, our analysis moves beyond existing literature focusing on innovation output levels to also examine how automation reshapes the *orientation* of innovation pursued by firms—whether firms deepen their expertise in existing technological domains or pivot toward entirely new fields, and the aggregate implication of such reorientation. In doing so, we offer novel insights into the reallocation of innovative effort under technological disruption and contribute to a broader understanding of firm dynamics in the age of intelligent machines. In particular, we highlight the complementarity between automation and AI, and document a pattern suggestive of an innovation "J-curve"—automation-induced innovation shifts to complementary domains such as AI incur high costs without immediate benefits, analogous to the productivity "J-curve" (e.g., Brynjolfsson, Rock, and Syverson 2021).

We start by extracting textual information from firms' 10-K filings to quantify their automation (interchangeably, robot) as the frequency of refined robot-related keywords normalized by total word count. This standardized metric for firm adoption and integration of robotic technology overcomes the limitations of industry-level proxy methods used in previous studies (e.g., M. B. Zhang 2019; Acemoglu and Restrepo 2020). Using an instrumental variable induced by the demographic changes

[1] As described for the background of the 2025 Nobel Prize in Economics,, the world has witnessed more economic growth than ever before over the past 200 years, the foundation of which is the constant flow of technological innovation.

in the European Union as in Acemoglu and Restrepo (2022), we study a plausibly causal link between automation exposure and innovation activities going forward. We find that firms with high robot exposure experience a notable shift in innovation focus from non-AI to AI-related technologies, a transition accompanied by a marked decline in patenting activity and a sharp rise in R&D expenditure. Upon closer scrutiny, the short-term reduction in innovation output is offset by greater generality and originality of innovation, underscoring the resource reallocation required for AI-driven innovation. Our paper therefore offers initial insights into how innovation begets innovation.

Specifically, we start by examining the general trends in firms' technological shift using two complementary metrics: (i) *Technology Similarity* captures continuity in a firm's patent portfolio composition by representing each portfolio as a vector in a space with basis dimensions corresponding to the firm's patent activity across categories. To track the evolution of a firm's technological focus, we construct annual patent distribution vectors over a five-year horizon and apply cosine similarity to assess its changes over time. (ii) *Semantic Similarity* leverages PatentBERT—a large language model specifically fine-tuned for generating semantic representations of patent texts—to detect subtle shifts in innovation themes, allowing for more precise year-over-year comparisons of innovation content beyond classification-based analyses.

Our empirical analysis reveals that heightened automation exposure drives a strategic reorientation of firms' innovation activities. Firms with greater robot exposure experience a decline in both technology similarity and semantic similarity, highlighting firms' adaptive recalibration in response to technological advancements. The findings remain robust across various alternative measures of automation exposure, including the industry-level metric proposed by Acemoglu and Restrepo (2019), as well as alternative instrumental variable approaches, such as the shift-share instruments discussed in Borusyak, Hull, and Jaravel (2022).

Next, we visualize the relative shifts in innovation direction across firms, finding that those with high automation increasingly patent in fields closely linked to AI. Building on this observation, we focus on AI as a key technological domain toward which innovation reorients in response to automation. The shift appears to be driven by two main factors: first, AI enhances robotic functionality by improving operational capabilities; second, AI benefits from the extensive data generated by robots, creating a self-reinforcing cycle of AI-driven advancements. These patterns provide preliminary evidence that

AI is a primary area of technological adaptation under increased automation.

To explore AI-focused innovation, we employ a keyword-plus-CPC approach to empirically identify AI-related patents. Specifically, we classify a patent as AI-related if its abstract contains at least one AI-specific keyword, as in Cockburn, Henderson, and Stern (2018), or if its primary and/or secondary Cooperative Patent Classification (CPC) codes match those identified by WIPO (2019) as corresponding to AI technologies. This approach addresses the limitations of the CPC system, which does not explicitly distinguish between AI and non-AI patents.

Beyond binary AI-versus-non-AI classification, we further invent two complementary metrics: AI closeness and distance. These measures quantitatively assess a patent's proximity to a conceptual "representative" AI patent, constructed by aggregating AI patents to mitigate the influence of outliers. Leveraging a patent-to-patent citation network, AI closeness measures the similarity in citation patterns between a given patent and the AI archetype, providing a more granular assessment of AI relevance, while AI distance captures the citation-network distance between a given patent and the "representative" AI patent. In addition to these relative measures, we also examine the absolute number of AI patents as a robustness check, evaluating whether firms' AI patenting activity increases in response to automation exposure.

Our findings reveal a clear and robust pattern: greater robot exposure is associated with a significant increase in patent closeness to AI-related technologies, alongside a corresponding reduction in patent distance from AI innovations. In addition, firms with high robot exposure are more likely to produce a greater number of AI-related patents in subsequent years, indicating a strategic shift towards AI-driven innovation. To assess the persistence of this innovation shift, we further analyze its dynamics over time. The results show that both AI closeness and the number of AI patents remain elevated even three years after initial robot exposure, indicating that the transition towards AI-focused innovation is not transitory.

Concerning the underlying mechanism driving this transition, our textual analysis reveals that firms with heightened robot exposure experience a surge in data-generation activities, pointing to a production externality. This is consistent with Acemoglu and Restrepo (2020), who find that automation primarily substitutes for low-skill labor, making tasks more standardized and codifiable—thus enabling the generation of richer, more structured data. Larger firms, in particular, show the greatest

data gains from robot exposure, consistent with their superior ability to scale non-rival data. This piece of empirical evidence suggests that firms strategically reorient their innovation efforts towards AI as a direct response to the integration of robotics, leveraging data-driven insights to accelerate AI advancements.

For robustness and more granular analyses, we use two alternative AI measures. The first is the USPTO AI dataset, which applies a machine learning model to classify U.S. patents and categorizes AI technologies into eight subclasses: machine learning, evolutionary algorithms, natural language processing, speech processing, vision processing, knowledge representation, planning and control algorithms, and hardware. The second is the AI patent dataset by H. Fang et al. (2025), who improve upon the USPTO classification by fine-tuning a domain-specific language model that achieves over 90% precision and recall. We find that firms with higher robot exposure exhibit innovation growth across all eight AI subclasses.

In addition to the direction of innovation, we explore how automation exposure influences general innovative activities, including both their inputs and outputs, which is unclear a priori. Concerning innovation outputs, we find a pattern suggestive of a J-curve–type dynamic, with a short- to medium-run decline followed by partial attenuation in the longer run in citation-weighted patent counts. Specifically, as robot exposure intensifies, the cumulative number of patents declines in the first two years, although the magnitude of the decline gradually diminishes over time. By the fifth year, the coefficient on robot exposure turns positive, and by the sixth year, the estimates suggest a moderate increase in the number of patents. This pattern suggests that firms initially experience a contraction in patent outputs as they adapt to automation, but this adjustment phase is followed by a recovery period in which innovation accelerates.

Concerning innovation inputs, we document a uptick in R&D investments, accompanied by a short-term decline in patent outputs following increased automation exposure (a pattern consistent with R&D's dual role in knowledge creation and absorptive capacity; see Griffith, Redding, and Van Reenen 2003). Specifically, our instrumental variable regression reveals a pronounced increase in R&D expenditure relative to total assets and sales, with a one-standard-deviation rise in robot exposure corresponding to 10.4% and 38.5% increases in the R&D-to-total assets and R&D-to-sales ratios, respectively. These shifts are statistically significant with substantial economic implications,

especially when contrasted with the mean R&D-to-total assets and R&D-to-sales ratios of 44.9% and 11.6%. The findings are also consistent with Babina et al. (2024), who demonstrate that AI investments stimulate firms' long-term innovation output by increasing R&D expenditure without generating immediate effects. Policymakers and firms therefore need to be patient with the long-term investments associated with such innovation shifts.

This puzzling gap between R&D spending and patenting may be attributed to the more foundational research with greater generality and originality, the high transition costs of specialized human or fixed capital investments required for new technology fields (McElheran et al. 2025), and the elevated marginal costs associated with these new research areas (Jones 2009). Specifically, the transition of firms to AI, particularly in cutting-edge machine learning or large models, is costly, but enables breakthroughs in fundamental science and technology (Jumper et al. 2021). Mezzanotti and Simcoe (2023) show that patenting is more closely tied to development than research. Therefore, if the increase in R&D is predominantly driven by more exploratory, original, and fundamental work—typically associated with the "R" in R&D—it is plausible that firms may be investing more in R&D while patenting fewer, albeit higher-quality, patents. Adopting the methodology in Hall, Jaffe, and Trajtenberg (2001), we indeed find that AI patents exhibit significantly higher levels of both generality and originality.

To further dissect the impact of transition costs on innovation shift, we delve into the potential mitigating effect of prior AI research experience. The premise is that familiarity with AI could ease the burdens of human capital acquisition during such a strategic pivot. Thus, we anticipate that firms with a rich background in AI research prior to 2004 would demonstrate a more pronounced shift towards AI-related innovations than their less experienced counterparts. Our analysis demonstrates empirically that firms with extensive AI research history indeed undergo greater innovation shift, as reflected in increased closeness to and decreased distance from AI-related patents.

Our investigation also touches upon the nuanced cost dynamics of AI-related innovation relative to non-AI fields, employing a series of metrics designed to gauge the marginal costs of innovation. This follows the approach of Jones (2009), who develops a theory to explore the implications of technology-specific human capital for technological progress. Our analysis reveals that AI innovation, compared to its non-AI counterparts, involves more team collaboration, greater labor input from growth-oriented inventors, and higher levels of inventor originality. In addition, AI patents tend to

generate higher value. These findings further corroborate the idea that automation exposure spurs firms to focus more on fundamental and large-scale innovations, which require greater inputs and take longer to yield observable outputs.

We demonstrate the robustness of our core findings through an array of additional tests involving alternative constructions of key variables, providing heuristic evidence on the validity of IV estimates (excluding common shocks and pre-existing trends, and including alternative IVs and falsification tests), controlling for the *Alice* decisions, and varying sample specifications. We also build a dynamic general equilibrium model to understand the underlying forces driving innovation shift amidst increasing automation and to rationalize the aforementioned empirical observations and our empirical construction of automation exposure. The model accommodates firms with varying degrees of automation exposure and allows for endogenous innovation choices. Firms engage in a range of tasks, with a subset automated and the remainder performed manually. Crucially, firms can innovate across these tasks, applying patented technologies to bolster task productivity. We show that the escalation in robot exposure, or more broadly automation, propels firms to recalibrate their innovation efforts towards domains that either enhance robot utilization or augment the efficiency of automated processes, given that innovations in tasks subject to automation act in synergy with existing robotic technologies.

Overall, our work adds to studies concerning multi-dimensional technological progress, particularly how market size and factor environments shape innovation incentives and technological direction. Prior research highlights the role of public innovation incentives (Finkelstein 2004), state-driven technology procurement (Beraja, Yang, and Yuchtman 2023), demographic fundamentals expanding demand (Acemoglu and Linn 2004; Costinot et al. 2019), labor-supply pressures (Acemoglu and Restrepo 2022), macro shocks reallocating inventive effort (Bloom, Davis, and Zhestkova 2021), social networks and homophily (Einiö, Feng, and Jaravel 2019), and innovation congestion and crowding externalities (Hopenhayn and Squintani 2021). Against this backdrop, we emphasize the role of recent automation adoption as a catalyst for shifting innovation toward AI. We document a technological complementarity between automation and AI, consistent with evidence of similar complementarities between data and AI (Eisfeldt, Schubert, and M. B. Zhang 2023), and and a pattern suggestive of an innovation "J-curve," in which legacy-field patenting declines before new innovative capacity materializes.

Our article also adds to studies resolving the puzzling decline in research productivity (e.g., Griliches 2009). In particular, Jones (2009) develops a theory of technology-specific knowledge accumulation, suggesting inventors' increasing demand for knowledge and learning to approach the innovation frontier. Bloom et al. (2020) employ industry-level case studies to similarly illustrate significant increases in research efforts despite stark declines in research productivity. We provide empirical evidence and microfoundation for these narratives. Innovation shifts prompted by automation—notably, towards AI-related fields—offer a tangible lens to view and understand part of the puzzle of declining research productivity. The shift in endogenous innovation, while strategically beneficial in the long run, can manifest itself in the seeming reduction of research productivity observed across industries.

Finally, our study enriches the discourse on the determinants of corporate innovation, which range from institutional and legal frameworks to market structures and individual corporate attributes. Recent studies have emphasized intellectual property regimes (L. H. Fang, Lerner, and C. Wu 2017), labor and bankruptcy laws (Acharya and Subramanian 2009; Acharya, Baghai, and Subramanian 2014), taxation (Akcigit et al. 2022), the evolution of the financial market (Hsu, Tian, and Xu 2014; Malamud and Zucchi 2019), market characteristics (Bloom, Schankerman, and Van Reenen 2013; Autor and Salomons 2018; De Bettignies et al. 2018), and specific corporate traits (Bena and K. Li 2014; J. Sunder, S. V. Sunder, and J. Zhang 2017; Frésard, Hoberg, and Phillips 2020) in shaping the innovation landscape. Our work adds a critical dimension by analyzing how technological progress directly influences firm innovation strategies, focusing on automation as a pivotal force reshaping industry-wide innovation patterns. While conventional analysis focuses on aggregate patent outputs, our novel metrics tracing the evolution of patent distributions across technological domains allow us to uncover nuanced patterns in how innovation shifts reallocate patenting activity across fields.

The remainder of the paper is structured as follows. Section 2 describes the data sources and empirical strategies. Section 3 presents the baseline results on how rising exposure to robotics drives shifts in firms' innovation direction. Section 4 explores the aggregate implications of these innovation shifts for the broader technological landscape. Section 5 reports robustness checks. Section 6 introduces a theoretical framework that rationalizes our findings. Finally, Section 7 concludes.

2 Data, Variables, and Empirical Strategy

2.1 Data Description

Our empirical analysis draws on several data sources, which we describe below.

Patents. Patent-related data are obtained from the United States Patent and Trademark Office (USPTO), which provides comprehensive information on all patents granted in the U.S. since 1976.[2] We use data on patent applications and grants to measure firm-level innovation output. Following Kerr (2008) and Kogan et al. (2017), we use the application year to construct annual patent counts, mitigating the distortions caused by the often-lengthy lag between application and grant. We link patent records to publicly listed firms using the methodology developed by Kogan et al. (2017), enabling consistent firm-level matching with financial variables.

Automation. As data on industrial robot installations are available exclusively at the industry level, prior studies estimate firm-level exposure to automation using industry-level data on robot installations from the International Federation of Robotics (IFR) combined with Compustat segment-level sales data (e.g., M. B. Zhang 2019; Acemoglu and Restrepo 2020). We similarly construct robotics exposure from two complementary sources. At the industry level, we use data from the IFR, which compiles annual records of industrial robot installations reported by leading global manufacturers. The IFR data cover approximately 90% of industrial robot installations across countries including the United States, Finland, Sweden, Denmark, Italy, and France. At the firm level, we extract relevant keywords and mentions of robotics usage from 10-K filings of U.S. publicly traded firms, to capture direct textual narratives of automation exposure.

Industry-level controls. We supplement our core data with industry-level information from the Bureau of Labor Statistics (BLS), including wages, employment, and price indices, as well as investment and production data from the Bureau of Economic Analysis (BEA). For instrumental variable construction, we use cross-country data from the EU KLEMS Growth and Productivity Accounts, which

[2] Data available at USPTO Patent Data.

provide industry-level measures of labor, capital, materials, energy, and services inputs across European economies. Unless otherwise noted, we classify industries using two-digit Standard Industrial Classification (SIC) codes.

Firm-level financials. We obtain firm-level financial and operational variables from the Compustat North America, including R&D intensity, total assets, debt ratios, firm age, return on assets (ROA), and asset tangibility. We also incorporate Compustat Segment data to extract detailed information on firms' industry and revenue composition.

Our sample covers the period from 2004, the first year when IFR provides data on industrial robot installations in the United States, to 2019, the latest year for which IFR data are available when we began our analysis. As is standard in the literature, we exclude firms in the financial and utility sectors due to their distinct regulatory and operational environments. We also drop foreign-incorporated entities and retain only firms that have applied for at least one patent during the sample period. Finally, we drop firm-year observations with missing key financial variables. Table A1 summarizes the composition of the final sample, which is restricted to publicly listed firms that can be linked to their SEC 10-K filings (via CIK identifiers).[3] The table reports the number of such firms, the annual share of firm-year observations with at least one patent, and the total number of patents. Notably, close to half of the firm-year observations have at least one patent, indicating a relatively innovation-intensive sample.

2.2 Automation and Robots

Measuring firm-level robot exposure. Prior studies such as M. B. Zhang (2019) and Acemoglu and Restrepo (2020) infer firm-level robot exposure from industry-level data on robot installations provided by the IFR. Robot exposure is defined as the number of industrial robots per thousand workers in each industry, weighted by the firm's sales composition across segments. The resulting firm-level measure is given by:

$$\text{Robot Exposure}_{it}^{\text{IFR}} = \sum_{k} \frac{\text{Sales}_{ikt}}{\text{Sales}_{it}} \times \frac{\text{Robots}_{kt}}{\text{Workers}_{kt}}, \tag{1}$$

[3]The construction of the text-based automation exposure measure using 10-K filings is described in the next section.

where Sales_{ikt} denotes the sales of firm i in industry segment k in year t, and Sales_{it} is the firm's total sales in the same year. The term $\frac{\text{Robots}_{kt}}{\text{Workers}_{kt}}$ captures the industry-level intensity of robot adoption in segment k at time t.

One potential concern is that firms may operate in few segments, leading to a limited within-industry variation. To address this concern, we conduct textual analysis on 10-K filings, leveraging the Generative Pre-trained Transformer (GPT) model. Specifically, we construct a firm-level robot exposure measure by identifying mentions of robot-related keywords in 10-K filings from 2004 to 2019. These filings provide comprehensive disclosures on firms' business activities and have been shown to be effective for measuring firm behaviors and characteristics through textual analysis, as demonstrated by prior research such as Hoberg and Maksimovic (2015).

We follow a three-step approach to measure firm-level robot exposure:

(i) Keyword selection and refinement: We first compile a comprehensive set of keywords based on prior literature (Acemoglu and Restrepo 2019; Acemoglu and Restrepo 2020; Acemoglu and Restrepo 2022; Graetz and Michaels 2018; Caselli and Manning 2019; Dixon, Hong, and L. Wu 2021; Koch, Manuylov, and Smolka 2021; Guerreiro, Rebelo, and Teles 2022) and further refine it using definitions from IFR reports.[4] In addition, we exclude terms related to the sale of industrial robots to ensure that our measure focuses on adoption and integration rather than commercial transactions. The finalized keyword list is presented in Table A2.

(ii) Contextual validation using GPT: Next, we extract sentences containing the identified keywords from 10-K filings and use OpenAI's *ChatGPT* (GPT–4)[5] to assess their relevance to robot adoption. Specifically, we design prompts that instruct ChatGPT to determine whether a sentence indicates the firm's actual and current use of robots, explicitly excluding forward-looking or planned adoption. Only sentences classified as relevant are retained for further analysis, ensuring that our study focuses on meaningful references from corporate disclosure.

(iii) Construction of the robot exposure measure: Finally, firm-level robot exposure is then

[4]For details, see IFR Industrial Robots.

[5]We use the OpenAI API with the following settings: `temperature` = 0, `top_p` = 1, `max_tokens` = 256, `presence_penalty` = 0, and `frequency_penalty` = 0. The system prompt is: *"You are an expert in industrial automation and SEC reporting."* The user prompt is: *"Given the sentence below extracted from a firm's 10-K filing, indicate whether it provides direct evidence that the firm has adopted or currently uses physical robots in its operations (excluding statements about future plans, intentions, or potential adoption). Answer #Yes or #No and briefly justify."*

measured based on the frequency of these keywords in 10-K filings, normalized by total word count to account for document length variations:

$$Robot\ Exposure_{it} = \sum_{w} \frac{Robot\ Word\ Count_{wit}}{Total\ Word\ Count_{it}}, \tag{2}$$

where Robot Word Count$_{wit}$ represents the occurrences of keyword w for firm i in year t, and Total Word Count$_{it}$ is the total number of words in the firm's annual report text.[6] The resulting measure is then standardized across all U.S. firms with available data to have a mean of 0 and a standard deviation of 1, facilitating cross-sectional and temporal interpretations.

Validity and robustness of robot exposure measure. To validate our robot exposure measure, we implement two empirical tests. First, we examine the correlation between our firm-level robot exposure measure (aggregated to the industry-year level) and actual industry-level robot adoption constructed by Acemoglu and Restrepo (2020). Figure I illustrates this relationship.[7] We find a positive correlation between our exposure measure and the observed robot-per-worker metric, providing some supporting evidence for the validity of our proxy as a measure of real-world robot adoption.

Second, to assess the robustness of our text-based construction, we evaluate the sensitivity of the measure to the choice of keywords using a leave-one-out procedure. This test iteratively excludes each keyword from the construction of the exposure score to evaluate whether any single term disproportionately influences the results. The measure remains stable across all iterations without significant deviations, thereby demonstrating robustness to keyword selection. To further ensure robustness, we explore several alternative constructions of the robot exposure measure. These results are presented

[6] Automation- and robotics-related disclosures are typically concentrated in Item 1 (Business) and Item 7 (MDA) of firms' annual reports. As a robustness check, we restrict the analysis to these sections and recompute keyword frequencies and total word counts accordingly.

[7] Industry-level robot intensity is calculated as the number of robots per worker, using robot installation data from the International Federation of Robotics (IFR) and employment data from the Bureau of Labor Statistics (BLS). Due to discrepancies in industry classification systems between the IFR and BLS datasets, we manually align 15 industries following the approach of Acemoglu and Restrepo (2020). These include broad sectors such as Agriculture, Forestry, and Fishing; Mining and Quarrying; Electricity, Gas, and Water Supply; Construction; Education; and R&D, along with nine manufacturing industries including Food and Beverages, Textiles, Wood and Furniture, Paper and Printing, Plastics and Chemicals, Glass and Ceramics, Basic Metals, Metal Machinery, Electrical/Electronics, and Automotive and Other Vehicles. Furthermore, because IFR reports robot installations jointly for the U.S., Canada, and Mexico prior to 2011, we follow Acemoglu and Restrepo (2020) in using combined North American data, as the U.S. accounts for over 90% of the region's industrial robot market.

and discussed in Section 5.2.1.

Beyond issues of validity and robustness, a further concern is whether robot adoption exhibits a lumpy investment pattern—characterized by large, discrete investments in a single year followed by little to no subsequent adoption. Such patterns stand in contrast to technologies like artificial intelligence, which are typically adopted more incrementally over time (Babina et al. 2024), and may complicate the identification of sustained versus one-off effects on firm-level innovation. Humlum (2019) documents such lumpiness among Danish firms.

To assess whether lumpy adoption patterns characterize the U.S. context, we examine the temporal dynamics of robot exposure across firms in our sample. Figure A1 shows the annual distribution of robot exposure among U.S. firms, revealing an increasing trend with moderate year-to-year variation. To further investigate the persistence of adoption, we identify the top decile of firms with the largest annual increases in robot exposure and track their exposure changes in the following year. Figure A2 indicates that firms experiencing greater increases in exposure in one year often continue to exhibit noticeable changes—both increases and decreases—in subsequent years.

Taken together, these results suggest that robot adoption, and consequently automation, among U.S. firms does not follow a predominantly lumpy pattern but rather reflects a more continuous and dynamic process. This observation supports the empirical identification of sustained effects of automation on firm behavior over time.

2.3 Identification Strategy and Instrumental Variable

We face two primary identification challenges. First, firms engaged in active innovation may be more likely to discuss robotics in their public disclosures, leading to measurement-induced reverse causality. Second, the adoption of robotics may be endogenous to broader structural or technological changes within firms or industries, thereby confounding causal inference.

Instrument construction. To mitigate these concerns and establish causal identification, we implement an instrumental variable (IV) strategy that builds on the approach introduced by Acemoglu and Restrepo (2020). Specifically, we exploit variation in demographic-driven robot adoption across advanced European economies—Denmark, Finland, France, Italy, and Sweden—where population

aging has long been a salient factor in shaping automation trends. These countries were among the earliest adopters of industrial robots, and demographic pressures have sustained high robot density since the 1990s. As Acemoglu and Restrepo (2020) argue, these early demographic shifts in Europe provide predictive power for subsequent patterns of robot diffusion in the United States.

We construct the instrument using industry-level robot usage data from IFR and employment data from the EU KLEMS Growth and Productivity Accounts. Due to differences in industry classification across these datasets, we manually reconcile and align 15 industries (SIC15), as detailed in Table A3. The instrumental variable for firm i in year t is defined as:

$$\text{Robot Exposure}_{it}^{\text{IV}} = \sum_{k} \frac{\text{Sales}_{ikt}}{\text{Sales}_{it}} \times \frac{\text{Robots}_{kt}^{EU}}{\text{Workers}_{kt}^{EU}}, \tag{3}$$

where Sales_{ikt} denotes the share of firm i's sales in industry k in year t, and Robots_{kt}^{EU} and Workers_{kt}^{EU} represent the number of industrial robots and workers, respectively, in industry k in year t across the selected European countries.

Validity and robustness of the instrumental variable. To assess the validity of our instrumental variable strategy, we conduct a series of tests[8]—to evaluate the key exclusion restriction, i.e., that robot adoption in European countries affects U.S. firm-level innovation outcomes only through its influence on U.S. automation exposure.

First, Acemoglu and Restrepo (2022) provide compelling evidence that demographic aging is a primary driver of robot adoption, and that aging trends vary substantially across countries. Their findings show that the advanced robotics adoption observed in countries such as Denmark, Finland, France, Italy, and Sweden is largely attributable to more rapid increases in the aging ratio relative to the United States. These demographic shifts are exogenous to U.S. firms and generate plausibly exogenous variation in international automation trends.

Second, we empirically rule out alternative channels through which European robot adoption might influence U.S. firm innovation. In particular, we test for confounding industry-level shocks—such as global price movements or industry-specific demand conditions—and the results provide limited

[8] See the regression results in, for example, Table I, as well as Section 5.1.

support for the view that such factors are driving our findings.

Third, we confirm the absence of differential pre-trends using the approach of Acemoglu and Restrepo (2020). By examining the relationship between robot exposure and firm innovation outcomes prior to 2004, we find no statistically or economically significant correlations. This result supports the identification assumption that post-2004 variation is not simply a continuation of prior trends. As an additional robustness check, we control for lagged values of the dependent variable from 1988 to 2003 to further mitigate concerns of omitted historical dynamics. We also implement a battery of formal IV diagnostic tests, including the weak identification test, underidentification test, and endogeneity test.

Finally, we assess the robustness of our findings through several alternative specifications. These include implementing a shift-share IV approach following Borusyak, Hull, and Jaravel (2022), introducing time-trend adjustments to account for differential automation trajectories, and expanding the set of EU countries used in constructing the instrument. Results across these specifications are generally similar, suggesting that our identification strategy is reasonable.

3 Robot Exposure and Innovation Shift

Our first objective is to assess the influence of robot exposure on the innovation trajectory of firms, explicitly investigating the extent to which increased exposure to robot catalyzes shifts in innovation focus towards AI-oriented domains. To this end, we first examine the broader trends in the technological shift among firms. We then focus on the impact of robot exposure on firms' propensity to pivot their innovation efforts towards AI-related fields. Finally, we explore one potential mechanism through which increased robot exposure drives a shift towards AI innovation.

3.1 General Innovation Shift

We construct vector representations of the patent portfolios of firms and use cosine similarity to quantify changes in innovation direction over time. We implement two complementary approaches to capture the various dimensions of technological content.

Technology similarity. We first measure innovation shift using the distribution of patents across technology classes, as classified by the USPTO's Cooperative Patent Classification (CPC) system. Among its hierarchical structure—30 two-digit, 126 three-digit, and 644 four-digit classes—we use the 126 three-digit level to balance granularity with interpretability.

Let N denote the number of technology classes. For firm i in year t, its patent distribution is represented by the vector $T_{it} = (T_{i1,t}, T_{i2,t}, \ldots, T_{iN,t})'$, where $T_{ij,t}$ is the share of patents filed in class j. To smooth annual fluctuations, we construct T_{it} using patents filed during the five-year window[9] from $t-4$ to t, and equally distribute a patent's weight across all assigned classes.

Following Bloom, Schankerman, and Van Reenen (2013), we compute technology similarity as the cosine of the angle between a firm's patent distribution vectors at years t and $t+5$:

$$\text{Tech Similarity}_{it} = \frac{T_{it}' T_{i,t+5}}{\sqrt{T_{it}' T_{it}} \sqrt{T_{i,t+5}' T_{i,t+5}}}. \tag{4}$$

This measure captures the persistence in a firm's technological focus. Higher similarity indicates a continuity in innovation domains, while lower similarity reflects a shift into new technological areas.

Semantic similarity. As an alternative to classification-based measures, we assess innovation similarity by analyzing the textual content of patent abstracts, which encapsulate the core conceptual contributions of each invention. Traditional approaches such as the bag-of-words (BoW) model quantify word frequency distributions but fail to capture semantic and syntactic relationships, as they treat words as independent and context-free. This limits their ability to represent the underlying meaning of technological content.

To address these limitations, we implement a transformer-based natural language processing approach using PatentBERT (Bekamiri, Hain, and Jurowetzki 2021), a BERT-based language model pre-trained and fine-tuned on a large corpus of patent documents. PatentBERT produces dense, 768-dimensional embeddings for each abstract, offering a high-fidelity representation of its semantic structure.

[9]The choice of a five-year window for patent-based proxies is common in the literature; see, for example, Rothaermel and Deeds (2004) and Hirshleifer, Hsu, and D. Li (2018).

We compute a firm-level semantic innovation profile by averaging the embeddings of all patents granted to firm i between $t-5$ and t, weighted by forward citations to account for heterogeneity in technological importance. Denoting this citation-weighted average embedding as T_{it}, we then measure semantic similarity as the cosine similarity between T_{it} and $T_{i,t+5}$, consistent with Equation (4). This metric captures the extent to which a firm's innovation content remains aligned over time: higher similarity indicates persistence in thematic focus, while lower values reflect a shift toward novel or exploratory technological domains.

Empirical specification. To test whether higher robot exposure prompts firms to redirect their innovation efforts towards different technological domains, we estimate the following:

$$y_{it} = \beta \times \textit{Robot Exposure}_{it} + \mathbf{X}'_{it}\gamma + \delta_i + \eta_t + \varepsilon_{it}, \tag{5}$$

where, y_{it} denotes the similarity of firm i between year t and year $t+5$. A lower value of y_{it} indicates a greater departure from the incumbent technological domains of the firm, consistent with a shift towards new areas of innovation. The key independent variable, $\textit{Robot Exposure}_{it}$, is measured as described in Section 2.2. The specification includes firm-fixed effects (δ_i) and year-fixed effects (η_t). Firm-fixed effects control for time-invariant, unobserved heterogeneity across firms, whereas year-fixed effects account for common temporal shocks, such as macroeconomic fluctuations or economy-wide technological trends. These controls ensure that identification relies on variations within or between firms over time. $\mathbf{X}_{it}$ includes standard control variables that vary over time, including the size of the firm (log of total assets), the leverage (debt ratio to total assets), the age of the firm (log), the profitability (return on assets, ROA) and the tangibility of assets (tangible assets as a proportion of total assets). The error term, ε_{it}, captures unobserved shocks. Standard errors are heteroskedasticity-robust and clustered at the firm level to account for within-firm serial correlation.

Results. Table I presents our main findings, indicating a shift in firms' innovation trajectories associated with increased exposure to automation. Columns (1) and (3) report the first-stage regression results, showing a significant positive association between the instrumental variable and robot exposure. The weak identification test yields an F-statistic of 24.00, well above the conventional threshold of 10,

alleviating concerns about weak instruments. Columns (2) and (4) report results using two measures of technological similarity. We find that greater exposure to automation is associated with a decline in both technological and semantic similarity relative to firms' prior patent portfolios. The effect is more pronounced for semantic similarity, with the coefficient in Column (4) statistically significant at the 5% level. Given that the mean technological similarity is 0.690, this represents a nontrivial reduction, about 63% of the mean, suggesting that firms facing greater automation pressures are significantly more likely to move away from their incumbent technological domains. Table A4 presents robustness results from specifications with and without controls, and reports the coefficients on the control variables. The estimates are broadly in line with the baseline findings.

3.2 AI-Oriented Innovation Shift

The initial findings indicate that firms with greater exposure to robots tend to adjust their innovation strategies, moving away from their established technological domains. Next, we explore the specific fields towards which these firms gravitate.

Innovation trajectory. To provide a clearer visualization of the evolution in innovation directions over time, we classify firms into high- and low-robot-exposure groups and compute the annual relative frequencies of n-grams in their patent abstracts. We then visualize the differences using a word cloud[10] (Figure II), which suggests that, relative to low-exposure firms, high-exposure firms place greater emphasis on software, electronics, machinery, and energy—domains that are often associated with AI-related applications.[11] Specifically, in 2010, n-grams such as "automated banking", "transmitting devices", "interface object", and "graphical program" emerge, which may reflect early signs of potential AI-related integration. By 2015, the prevalence of terms like "turbine engine", "gas turbine", "downhole tool", and "drill bit" indicates a continued emphasis on mechanical and energy

[10] Specifically, we first classify firms into high and low robot exposure groups. We then preprocess each patent abstract by converting text to lowercase, removing HTML entities, and eliminating non-alphabetic characters. After excluding stop words, we construct contiguous 2-grams (and higher-order n-grams) and compute their annual frequencies separately for the high- and low-exposure groups. Within each group, frequencies are normalized relative to the baseline year of 2004. We then take the ratio of these normalized frequencies between the two groups and apply a natural log transformation to obtain group-level log-ratios of n-gram prominence. The top-weighted log-ratios are visualized using a circular-masked word cloud.

[11] Table A5 presents the 20 n-grams with the highest relative frequency, corresponding to the results in Figure II.

technologies, some of which may incorporate AI-enabled monitoring, predictive maintenance, or automated control. By 2019, the presence of terms such as "transmission medium", "syntax elements", "electromagnetic wave", and "communications described" is consistent with a growing overlap between AI and communication or information-transmission technologies, where AI may play a role in signal optimization, real-time decision-making, or large-scale data processing. This trajectory is suggestive of a broader diffusion of AI-related techniques across a range of industrial domains.

Our analysis focuses on AI as a key technological domain toward which firms may redirect their innovation efforts in response to rising robot exposure. Section 6 provides a model to formalize this mechanism. The underlying intuition can be summarized as follows. From the return side, as automation expands the share of tasks performed by machines, it increases the marginal return to innovation in automated processes, thereby raising the relative payoff of technologies—such as AI—that enhance the productivity and functionality of robotic systems. From the cost side, robotic production generates structured operational data that can be used to train and improve AI models, which may reduce the effective cost of AI innovation through learning and data-driven feedback. From a resource reallocation perspective, given convex R&D costs, firms may reallocate innovation effort toward domains with relatively higher returns and lower effective costs, leading to a gradual shift away from non-automated technologies toward AI-related innovation. Taken together, these forces suggest a mutually reinforcing relationship between automation and AI innovation, contributing to the increasing integration of AI into firms' innovation activities.

Identifying AI patents. A key challenge in identifying AI patents arises from the absence of explicitly defined AI-specific classifications within the CPC system, mainly due to the broad and rapidly evolving scope of AI. Previous studies have addressed this challenge primarily through two approaches: (1) combining keyword searches in patent abstracts with CPC codes for identification (Cockburn, Henderson, and Stern 2018; WIPO 2019); and (2) employing machine learning methods trained on AI patent seed sets (Giczy, Pairolero, and Toole 2022; Miric, Jia, and Huang 2023; H. Fang et al. 2025). Following this literature, we adopt the keyword-plus-CPC approach as our primary methodology for identifying AI patents, supplemented by the machine learning approach as a robustness check. Specifically, we classify a patent as an AI patent if its abstract contains at least

one keyword from a predefined set listed in Table A6, which is adapted from Cockburn, Henderson, and Stern (2018), or if its primary and/or secondary CPC codes match those used by WIPO (2019) to identify AI-related technologies.[12]

To validate our keyword-based approach for identifying AI patents, we conduct robustness checks using two alternative datasets. The first is the USPTO AI patent dataset developed by Giczy, Pairolero, and Toole (2022). While the USPTO dataset has the limitation of relatively modest classification accuracy, it provides a valuable basis for more granular analyses due to its detailed categorization of patents into AI subclasses. Specifically, the USPTO dataset is constructed by initially creating a manually reviewed seed dataset to differentiate AI-related patents from non-AI patents. Subsequently, a machine learning algorithm— the Long Short-Term Memory (LSTM) model—is applied to classify the entire universe of U.S. patents. The second is the AI patent dataset constructed by H. Fang et al. (2025), who further improve upon the USPTO classification by fine-tuning a domain-specific language model to achieve higher precision and recall. One concern is that similar technologies may be described with fewer AI terms earlier and more explicit AI terminology later; however, time-varying seed keywords in both datasets help mitigate terminology-driven misclassification.

3.3 AI Closeness and Distance

After identifying AI patents, this section introduces two measures to quantify the proximity between any given patent and AI patents.

AI closeness. We measure a patent's thematic proximity to artificial intelligence using a citation-based metric termed *AI Closeness*. This measure compares the distribution of citations made by a given patent to those made by a composite "representative" AI entity.

For each year t, we aggregate all AI patents filed at t into a single representative entity and construct its citation distribution across USPTO technology classes:

$$\text{Cites}_{AI,t} = \left(\text{Cites}_{AI,1t}, \ldots, \text{Cites}_{AI,Nt}\right),$$

[12] Our sample period largely predates the widespread adoption of large language models (e.g., GPT), and thus primarily captures AI technologies in the pre-generative AI era.

where $\text{Cites}_{AI,nt}$ denotes the number of backward citations made by AI patents to technology class n. We normalize this to obtain a distribution vector:

$$C_{AI,t} = \left(C_{AI,1t}, \ldots, C_{AI,Nt}\right), \quad \text{with} \quad C_{AI,nt} = \frac{\text{Cites}_{AI,nt}}{\sum_n \text{Cites}_{AI,nt}}.$$

For any patent p filed in year t, we construct a comparable citation distribution vector C_{pt}. The *AI Closeness* of patent p is then defined as the cosine similarity between C_{pt} and $C_{AI,t}$:

$$\text{AI Closeness}_{pt} = \frac{C'_{AI,t} C_{pt}}{\|C_{AI,t}\|_2 \|C_{pt}\|_2}, \tag{6}$$

where $\|\cdot\|_2$ denotes the Euclidean (L_2) norm.

To validate the measure that AI patents exhibit a higher *AI Closeness* than non-AI patents, we perform a holdout exercise: we randomly reclassify 10% of AI patents as pseudo non-AI and retain the remaining 90% as the reference group. The pseudo non-AI patents exhibit significantly higher AI Closeness scores than actual non-AI patents, with a mean difference of 0.118. Moreover, the distribution of AI Closeness for these pseudo non-AI patents is more right-skewed than that of true non-AI patents, suggesting that our metric captures proximity to AI innovations. (Figure A3).

AI distance. To complement the citation-based closeness measure, we also define a structural proximity metric, *AI Distance*, based on shortest-path length in the citation network.

we construct an annual undirected patent-to-patent citation graph, where two patents p and p' are considered connected if a citation exists between them, regardless of direction. The *AI Distance* of a patent is defined as the minimum number of citation links required to connect it to the representative AI entity. A direct citation implies a distance of one; greater distances indicate increasingly indirect connections to the AI representative. By definition, all AI patents have a distance of zero.

Results. We estimate the relationship between robot exposure and firms' orientation toward AI innovation using the specification described in Equation (5). The dependent variable y_{it} is defined as either the average AI Closeness or the logarithm of AI Distance (plus one), averaged across all patents granted to firm i in year t.

Table II reports the results from two-stage least squares (2SLS) estimation. Panel A presents the main firm-level analysis. Columns (1) and (3) show the first-stage regressions, while Columns (2) and (4) report the second-stage estimates. The results indicate that greater exposure to robotics is associated with stronger AI alignment. A one standard deviation increase in robot exposure leads to a 0.004 increase in AI Closeness and a 0.736 reduction in AI Distance, both statistically significant at the 5% level. Relative to sample means of 0.002 and 2.779, these effects are economically meaningful, suggesting that automation plays a meaningful role in shaping the technological direction of innovation.

Panel B conducts robustness checks at the patent-firm-year level. Here, the dependent variable is defined at the individual patent level—either AI Closeness or log AI Distance (plus one)—and regressed on firm-year robot exposure, controlling for firm-level characteristics. The results are consistent: patents from firms with higher robot exposure are significantly more semantically proximate to AI and structurally closer within the AI citation network. Both coefficients are statistically significant at the 1% level.

Panel C examines dynamic effects. Columns (1)–(3) show that the positive impact of robot exposure on AI Closeness persists for up to three years, indicating that firms continue to deepen their engagement with AI-related innovation well beyond initial exposure. Columns (4)–(6) provide evidence that the effect of robot exposure on AI Distance remains statistically significant in the following year at the 10% level.

Panel D shifts the focus to innovation intensity by reporting the effects of robot exposure on the number of AI patents granted in the current year and in the one-, two-, and three-year periods thereafter. A one standard deviation increase in robot exposure is associated with an additional 0.211 AI patents in the current year, and this effect persists over the following two years. These findings underscore both the directional and intensive margins through which automation drives AI-oriented innovation at the firm level.

Alternative measure of AI patents. To further validate our findings and provide finer-grained insights, we employ two complementary AI patent datasets described in Section 3.2. The first is the USPTO AI patent dataset, which classifies AI-related patents into eight distinct technological sub-

classes: machine learning (ML), evolutionary algorithms (EVO), natural language processing (NLP), speech processing (SPEECH), vision processing (VISION), knowledge representation (KR), planning and control algorithms (PLAN), and AI hardware (HW). This classification enables us to investigate how firm innovation responses to automation differentially across specific areas within the AI landscape.

The second dataset is the AI patent dataset constructed by H. Fang et al. (2025), who document potential misclassification in the USPTO's AI patent tagging system and propose an improved classification approach using a fine-tuned AIPatentBERT model. Their algorithm achieves precision and recall rates exceeding 90%, representing a significant advancement over the existing USPTO classifier. In this paper, we directly leverage their refined AI patent data.

To address concerns regarding subclass-level classification accuracy, we aggregate all patents within each AI subclass into an annual "representative" patent. For each subclass and year, we then compute AI Closeness and AI Distance using the methodology outlined above.

Table III presents the results of the subclass-level analysis using the AI subclass classification proposed by H. Fang et al. (2025), whereas Table A7 relies on the USPTO's original classification. In Table III, Panel A shows that firms with greater robot exposure exhibit higher AI Closeness across all eight AI subclasses. The estimated coefficients are positive and statistically significant at the 5% level, suggesting that automation is associated with a reorientation of innovation toward AI-related domains. Panel B reports the corresponding estimates for AI Distance. While most coefficients are negative, they are not statistically significant. A plausible explanation lies in the construction of the AI Distance metric. Specifically, AI Distance is defined as the shortest path from a given patent to any AI patent in the citation network—a minimum statistic that is highly sensitive to classification error. If the underlying AI patent identification is noisy or incomplete, the measurement error introduced is likely to attenuate the estimated effect, biasing coefficients toward zero. As a result, the lack of statistical significance in Panel B may reflect imprecision in the distance measure rather than the absence of a true relationship.

Panels A and B of Table A7, which rely on the USPTO-based AI classification, yield results that are qualitatively similar to those in Table III. While the USPTO measure is noisier, the consistency in directional patterns reinforces the robustness of our findings.

Taken together, these results suggest that robot exposure not only reorients the direction of firm innovation but also increases engagement with the technological frontier of AI across a broad spectrum of specialized domains. This highlights automation's role as a catalyst for frontier innovation in the era of intelligent technologies.

3.4 Mechanism: Rising Data Generation

This section examines a potential mechanism through which increased robot exposure fosters AI innovation: enhanced access to data. Building on the findings of Beraja, Yang, and Yuchtman (2023), which demonstrate that government procurements in China significantly stimulate AI innovation by granting access to proprietary data, we investigate whether firms with higher robot exposure similarly gain access to richer data resources. This improved data availability may accelerate AI advances, driving breakthroughs in AI technologies and applications. In turn, these innovations can further enhance and optimize the development and performance of robotics systems, creating a mutually strengthening cycle of technological progress.

Amazon and Kiva: a case study. A notable example is Amazon's acquisition of Kiva Systems in March 2012. As one of the world's largest e-Commerce and cloud computing companies, Amazon develops advanced AI systems to optimize its operations. Kiva Systems, on the other hand, specializes in mobile robotics for warehouse automation. By incorporating Kiva's robotics technology, Amazon not only improves warehouse efficiency but also gains access to vast data streams, fostering new opportunities for AI-driven innovation. In turn, this facilitates the ongoing enhancement and evolution of Kiva's robotic systems, further optimizing their efficiency, adaptability, and overall performance in warehouse automation.

To be more precise, Kiva's robotic systems generate data across multiple dimensions that prove crucial for Amazon's AI-driven innovations. Operational metrics (e.g., position, speed, and trajectory) enable the optimization of robot movements, reducing errors, and boosting efficiency. Performance indicators (e.g., cycle times and accuracy rates) offer system-wide insights for refining logistics strategies, while environmental data (e.g., temperature and humidity) facilitate optimal working conditions. Energy usage records inform cost-effective resource allocation, and maintenance logs (including fault

tracking and predictive maintenance schedules) minimize downtime and extend equipment lifespans. Collectively, these data streams fuel Amazon's AI capabilities, driving continuous improvement in robotics and warehouse operations.

Measuring data generation. To examine whether increased robot exposure leads to increased data availability, we must quantify the extent of data accessible to firms, which is not straightforward to measure. Following an approach similar to that used to measure robot exposure, we use the frequency of terms related to data generation in 10-K filings as a proxy for data access. The keywords, detailed in Table A8, are derived from a broad body of literature (e.g., Abis and Veldkamp 2024; Bajari et al. 2019; Brynjolfsson and McElheran 2016; L. Wu, Hitt, and Lou 2020; C. Zhu 2019).

Our methodology consists of keyword selection, contextual validation using GPT to identify and retain statements pertaining to data generation (as opposed to data use), and computing keyword frequency in firms' 10-K annual reports, which is then normalized by total word count to account for variations in document length. Formally, we define the data generation measure as:

$$Data\ Generation_{it} = \sum_{w} \frac{Data\ Word\ Count_{wit}}{Total\ Word\ Count_{it}}, \tag{7}$$

where $Data\ Word\ Count_{wit}$ represents the count of data-related keywords w appearing in firm i's 10-K filing during year t and $Total\ Word\ Count_{it}$ is the total number of words in the corresponding report. We standardize this measure to mean zero and unit variance for comparability.

Results. Table IV presents evidence that firms with greater robot exposure experience a significant increase in data generation. Columns (1) and (2) report the two-stage estimates of the average effect of robot exposure on data generation. The results indicate that a one-standard-deviation increase in robot exposure is associated with an increase in data generation of approximately 1.313, which is statistically significant at the 1% level. Columns (3)-(5) further examine heterogeneity by firm size. Specifically, we categorize firms into two groups based on sales: small firms (bottom 75%) and large firms (top 25%). The results reveal that the interaction terms between robot exposure and the large-firm indicator are positive and marginally significant. These findings provide suggestive evidence that larger firms may derive greater benefits, exhibiting a relatively larger increase in data generation as

robot exposure increases. This pattern is broadly consistent with the non-rival nature of data, whereby larger firms may be better positioned to leverage data-driven advancements.

4 Aggregate Innovation

Although we establish that increasing automation prompts firms to reorient their innovation focus towards AI-related fields, the overall impact of automation on innovation outputs remains uncertain. Theoretically, automation could either stimulate or constrain a firm's innovation activities, depending on factors such as the transition costs to new technological domains, variations in research efficiency and complexity, and the levels of R&D investment. This section aims to investigate the effect of robot exposure on firms' innovation productivity.

4.1 Innovation Outputs

To evaluate the effect of robot exposure on innovation outputs, we examine two key metrics: the quantity of innovation, measured by patent counts, and the quality of innovation, proxied by citation-weighted patent value. Recognizing that innovation responses may be delayed, we analyze cumulative outcomes over multiple future years.

Empirical specification. Following Cohn, Liu, and Wardlaw (2022), we model patent generation as a time-varying Poisson process, appropriate for count data that are non-negative and sparsely distributed.[13] The expected number of patents filed by firm i from t to year $t+j$ is specified as:

$$\mathbb{E}(y_{i,t+j} \mid \mathbf{X}_{it}) = \exp\left(\beta \times \text{Robot Exposure}_{it} + \mathbf{X}'_{it}\gamma + \delta_i + \eta_t\right), \tag{8}$$

where δ_i and η_t are firm and year fixed effects, and $\mathbf{X}_{it}$ includes time-varying firm controls. The dependent variable is adjusted for patent quality by computing forward five-year citations and normalizing by the average citation count within each three-digit CPC class to adjust for field-specific

[13] Although we adjust the dependent variable using citation-based weights, it remains non-negative and skewed. Consistent with Gourieroux, Monfort, and Trognon (1984) and Silva and Tenreyro (2006), we adopt a Poisson regression model with conditional expectation $\mathbb{E}(y \mid X) = \exp(X\beta)$.

citation inflation, following Acemoglu, Akcigit, and Kerr (2016).

To assess patent value, we use a linear specification:

$$y_{i,t+j} = \beta \times \text{Robot Exposure}_{it} + \mathbf{X}'_{it}\gamma + \delta_i + \eta_t + \varepsilon_{it}, \tag{9}$$

where $y_{i,t+j}$ denotes the cumulative citation-weighted value of patents filed by firm i between years t and $t+j$, as defined by the methodology in Kogan et al. (2017).

Results. Table V reports the cumulative effects of robot exposure on patent outcomes. We find that robot exposure is associated with a statistically significant decline in the number of patents over the first two years, significant at the 5% and 10% levels, respectively. The magnitude of this decline diminishes over time. From three to five years after exposure, the estimated coefficients are statistically indistinguishable from zero, while the point estimates gradually increase and turn positive in the fifth year. By the sixth year, the coefficient becomes positive and marginally significant at the 10% level, suggesting a pattern broadly consistent with a potential "J-curve."

To address potential survivorship bias, Table A10 presents robustness checks. Panel A reports results based on a balanced panel, while Panel B presents estimates from an IV specification. The balanced panel results indicate that the cumulative number of patents begins to increase from the third period onward, with coefficients significant at the 5% level initially and remaining positive at the 10% level in subsequent periods. In contrast, the IV results suggest a modest decline in cumulative patenting in the earlier periods.

Combining these results with Table II, which shows a significant and persistent increase in AI-related patents, the observed pattern in total patent counts appears to reflect a shift in innovation composition. The initial decline is consistent with a reduction in non-AI patents, while the subsequent recovery appears to be gradually offset by increased AI-related patenting, consistent with a reallocation of innovation resources toward AI.

The pattern for patent value is more muted. While cumulative patent value declines significantly during the first four years—relative to a sample mean of 1.218—these effects dissipate thereafter, with coefficients becoming statistically insignificant beyond year four.

To further explore heterogeneous effects, we classify patents by two-digit CPC codes and interact

these categories with robot exposure in Equation (8). Figure A4 identifies the ten CPC categories most negatively affected, as well as those experiencing increases. The results show that most technological areas see a reduction in patenting activity. However, a subset of growing categories is closely aligned with AI-related technologies, including mechanical components, amusement devices, and checking instruments. As CPC codes do not explicitly identify AI patents, our dedicated AI classification remains essential for precision.

Taken together, these findings suggest that the impact of robot exposure on innovation is dynamic and time-dependent. While automation initially suppresses patent output and value, these effects partially reverse over time. This pattern aligns with Babina et al. (2024), who document that AI-driven innovation unfolds gradually, with long-run benefits that are not immediately observable. Our results highlight the potential for automation to induce temporary disruption followed by a strategic pivot toward higher-value, frontier innovation.

4.2 Innovation Inputs

We next examine how automation influences firms' innovation inputs by analyzing changes in their R&D expenditures. This analysis complements our earlier findings on innovation outputs and provides insight into potential reallocation dynamics that are suggestive of a J-curve–type pattern.

Empirical specification. Specifically, we analyze the impact of robot exposure on R&D intensity using the following estimation model:

$$y_{it} = \beta_1 \times \textit{Robot Exposure}_{it} + \beta_2 \times \textit{R\&D Dummy}_{it} + X'_{it}\gamma + \delta_i + \eta_t + \varepsilon_{it}, \tag{10}$$

where y_{it} represents firm i's R&D intensity, measured as R&D expenditure scaled by total assets or sales in year t. Given the incomplete disclosure of R&D data in Compustat, missing values are set to zero. To account for this imputation, we include a dummy variable, $\textit{R\&D Dummy}_{it}$, which equals one if firm i reports positive R&D expenditure in year t, and zero otherwise.

Results. Table VI reports the two-stage estimation results for the impact of robot exposure on firms' R&D expenditure, normalized by total assets and total sales. A one-standard-deviation increase in robot exposure leads to increases of 0.104 in the R&D-to-assets ratio and 0.385 in the R&D-to-sales ratio, significant at the 10% and 5% levels, respectively. Given mean values of 0.449 and 0.221, these effects are economically significant. The observed rise in R&D expenditure, coupled with the "J-curve" pattern in patent innovation linked to increased robot exposure, suggests a potential reallocation of innovation efforts, which we explore further in the subsequent subsections. [14]

4.3 Mechanism: Generality and Originality of Innovation

A plausible explanation for the observed innovation "J-curve" is that the shift towards AI-related innovation redistributes resources between research and development. Specifically, firms may allocate greater emphasis to fundamental research, leading to a temporary decline in patent quantity while enhancing patent generality and originality.

The distinction between research and development lies in their respective roles in knowledge creation. According to the National Science Foundation (NSF), research is the planned, systematic pursuit of new knowledge or understanding, whereas development involves the application of research and practical experience to create or improve goods, services, or processes. Mezzanotti and Simcoe (2023) show that patenting is more closely tied to development than to research. Thus, if the observed increase in R&D expenditure is primarily driven by exploratory, original, and foundational research—which typically falls under the research component of R&D—it is plausible that firms invest more in R&D while producing fewer but higher-quality patents with greater technological significance. In our context, firms' transition towards AI innovation often necessitates breakthroughs in foundational science and technology.[15] As a result, AI-driven patent innovation compels firms to revisit fundamental challenges, striving for original breakthroughs in AI theories and methodologies

[14] Chuang et al. (2024) show that omitting firm fixed effects helps explain the "missing link" between R&D input and patent output. In our analysis, we also find that excluding firm fixed effects yields a significantly positive relationship between robot exposure and both R&D intensity and patent counts. However, given the potential correlation between firm-level unobservables and our key explanatory variables, we retain firm fixed effects in our main specification to mitigate omitted variable bias and ensure more credible identification.

[15] Advancements in AI rely on profound understanding of underlying theories, algorithms, and hardware, as well as the interactions among them. To maintain a competitive edge in AI, firms must develop novel solutions, pushing the boundaries of algorithms, data structures, and processing methods.

that surpass existing technological constraints.

Measuring fundamental research. To empirically verify whether firms with greater automation exposure focus more on fundamental research, we adopt the approach outlined by Hall, Jaffe, and Trajtenberg (2001) to measure a patent's generality and originality:

$$Generality_{pt} = 1 - \sum_n s_{pnt}^2, \; Originality_{pt} = 1 - \sum_n q_{pnt}^2, \tag{11}$$

where s_{pnt} represents the share of citations that patent p filed in year t receives from technology class n within five years of its issuance, while q_{pnt} denotes the share of citations patent p makes to technology class n, with technology classes defined at the 3-digit CPC level. Intuitively, generality captures the extent to which a patent is cited across a broad range of technology classes, reflecting its impact beyond its specific domain. In contrast, originality measures how broadly a patent draws from diverse technological fields, indicating the multidisciplinary nature of its knowledge base. One potential concern with citation-based measures is that patents tend to receive citations from a wider array of technology classes over time. To address this, we measure the generality of a patent based on citations received within the five years following its filing year t, thus mitigating biases stemming from long-term citation trends.[16]

Results. Table VII presents the results. Panel A presents firm-year-level regressions, where the generality and originality are aggregated across all patents filed by a firm in a given year. The estimates show that higher robot exposure is associated with increases in both generality and originality of innovation. Specifically, a one-standard-deviation increase in robot exposure is associated with an approximately 0.135 increase in generality and a 0.028 increase in originality. These effects are statistically significant at the 10% and 5% levels, respectively, and are economically meaningful given mean values of 0.043 and 0.007. As a robustness check, Panel B presents regressions at the patent-firm-year level. The results remain consistent with the main findings and exhibit even stronger statistical significance, with all estimates significant at the 1% level.

These findings suggest that during the transition towards AI-driven innovation, firms place greater

[16]The generality measure excludes patents with no forward citation within the five-year observation window.

emphasis on foundational research characterized by higher generality and originality. This shift leads to an increase in R&D expenditure but a decline in patent output. Moreover, our findings provide new insights into the evolution of the United States scientific-industrial complex in the era of AI innovation. As highlighted in studies such as Arora, Belenzon, and Patacconi (2018), a key transformation since the 1980s has been the reallocation of resources and attention by leading firms away from exploratory scientific research towards more commercially oriented projects. However, our research identifies an opposing trend, wherein the shift towards foundational AI innovation can, in turn, stimulate more exploratory scientific research.

4.4 Mechanism: Standing on Jones' Shoulders

The second possible explanation for the innovation "J-curve" is the investment in human capital accumulation required for AI-related innovations. As Isaac Newton famously stated, "Innovation is just standing on the shoulders of giants". Learning from past advancements can significantly reduce the transition costs associated with human capital-specific innovation (Bloom, Schankerman, and Van Reenen 2013; Cui, Lu, and W. Zhu 2025). To capture this effect, we use firms' AI-related innovation experience prior to 2004 (the initial year of our sample) as a proxy for such transition costs. Firms with a greater prior experience in AI are likely to have developed specialized expertise and accumulated knowledge, which may facilitate a smoother and more efficient transition to AI-driven innovation. In contrast, firms with limited prior AI exposure may face higher barriers in adapting to AI-intensive research, as they need to invest more heavily in acquiring specialized talent and developing new capabilities.

Empirical specification. We measure a firm's AI research experience based on the number of AI patents granted between 1993 and 2003. Firms are classified as having AI experience if their AI patent count exceeds the industry median; otherwise, they are considered to have less experience. We then test whether firms with prior AI experience exhibit a more significant shift in research focus using the

following specification:

$$y_{it} = \beta_1 \times Dummy(AI\ Experience)_{i,03} \times Robot\ Exposure_{it} + \beta_2 \times Robot\ Exposure_{it} + \mathbf{X}'_{it}\gamma + \delta_i + \eta_t + \varepsilon_{it}, \tag{12}$$

where y_{it} represents the aggregated firm-level AI closeness or distance, calculated based on all patents of firm i in year t. $Dummy(AI\ Experience)_{i,03}$ is an indicator equal to 1 if firm i's AI experience prior to 2004 exceeds the industry median, and 0 otherwise. If prior research experience facilitates the transition towards AI-related innovations, we expect firms with greater AI research experience to exhibit a more pronounced shift towards AI fields in response to rising automation compared to their less experienced counterparts.

Results. Table VIII presents the heterogeneous effects of rising robot exposure on firms' shift toward AI innovation at the firm level, while Table A9 reports patent–firm–year–level estimates as a robustness check. Together, the results highlight two key insights. First, increased robot exposure is associated with a significant increase in AI closeness and a reduction in AI distance for both groups of firms. Second, the magnitude of this shift is nearly twice as large for firms with higher prior AI experience at the patent-firm-year level, suggesting that firms with a stronger AI foundation are better positioned to integrate automation-driven advancements into their innovation strategies. These findings highlight shift costs in shaping firms' adaptive responses to automation, reinforcing the importance of accumulated knowledge in facilitating technological transitions.

4.5 Mechanism: The Relative Cost of AI-Driven Innovation

Beyond the high transition costs, the rising relative costs of AI-driven innovation may also explain the increase in R&D expenditure and the initial decline in patent output—an adjustment burden also noted by Hill et al. (2025) in analyzing researchers' directional shifts. Building on the framework of Jones (2009), we examine the incremental costs of AI patents relative to non-AI patents across multiple dimensions using several constructed measures of innovation costs.

Measuring relative costs. **(i) Team size.** When innovation requires the accumulation of knowledge across broader technological domains, it is more likely to be undertaken by larger teams of inventors with complementary expertise. Accordingly, we measure a patent's team size as the number of inventors listed on the patent. A larger team size may indicate the increasing complexity of AI-related innovation and the necessity for interdisciplinary collaboration.

(ii) Labor input. Team size alone may not fully capture the actual labor input associated with a patent, as inventors often contribute to multiple projects simultaneously, potentially leading to double counting. To address this concern, we construct an adjusted measure of labor input that accounts for an inventor's participation across multiple patents. Specifically, if an inventor contributes to multiple patents filed within the same year, their labor input is evenly distributed across those patents.

(iii) Inventor originality. Hall, Jaffe, and Trajtenberg (2001) suggest that an inventor's originality is positively correlated with innovation costs. The underlying rationale is that generating highly original innovations necessitates searching for and integrating knowledge from a broader range of technological fields, thereby increasing the complexity and resource demands of the innovation process. We quantify an inventor's originality using the following measure:

$$Inventor\ Originality_{mt} = 1 - \sum_{n} s_{mnt}^2, \tag{13}$$

where s_{mnt} represents the fraction of patents by inventor m that fall into technology class n over the past five years. For each patent, we compute its inventor originality as the average originality score of all inventors associated with the patent.

(iv) Patent value. Patent value serves as an additional proxy for innovation costs, as higher-value patents generally require greater investments in research and development. A patent with significant economic impact is often the result of extensive exploratory work, considerable resource allocation, and rigorous development efforts. We measure patent value following the approach proposed by Kogan et al. (2017), which quantifies a patent's economic significance based on its market impact and citation.

Results. Figure III presents a comparative analysis of AI and non-AI patents across key cost-related dimensions. Subfigure (a) shows that AI patents tend to involve larger inventor teams than non-AI

patents, with the gap widening notably after 2015—coinciding with the growing prominence of AI-driven innovation. Subfigure (b) complements this pattern by documenting higher average labor input for AI patents, highlighting their greater complexity and resource intensity. Subfigure (c) illustrates differences in inventor originality, with AI patents typically associated with broader technological expertise. Although AI patents have historically scored higher on this measure, the gap has diminished in recent years, potentially reflecting the widespread diffusion of AI across diverse fields. Subfigure (d) highlights the value premium of AI patents over non-AI patents. The widening gap after 2011 suggests that AI-related inventions increasingly generate patents with greater technological and economic impact. To further validate these findings, we employ a regression approach with firm-year-fixed effects, and report the findings in Table A11. Together, these findings highlight the unique cost structure of AI innovations and their growing role in the technological landscape.

5 Robustness Checks

Next, we conduct several robustness checks to validate our core findings on innovation shift and innovation productivity, including tests of IV validity, robustness of text-based robot exposure measures, the effects of the Supreme Court's *Alice* decision, and alternative sample constructions.

5.1 Assessing the Validity of IV Estimates

5.1.1 Common shocks and pre-trend tests

The validity of using robot usage per worker in the EU countries as an instrumental variable may be subject to several potential challenges. One concern is that common demand shocks could simultaneously influence both EU robot investment and the innovation activities of US firms, leading to endogeneity issues. Another potential challenge is the presence of a pre-existing trend in the innovation direction of US firms prior to the start of the sample period, raising the possibility that the observed effects are not solely driven by automation exposure. To examine the robustness of our findings, we implement two robustness checks to address these concerns.

(i) Common shock effects. To account for potential confounding shocks, we examine whether

industry-level price dynamics influence our results, as fluctuations in output prices may affect firms' automation and innovation strategies. Specifically, we incorporate the industry-level Producer Price Index (PPI) into the regression framework and also control for industry-year fixed effects to capture any time-varying, sector-specific factors.

The results, reported in Table A12, show that the estimated coefficients remain largely stable after including these controls, indicating that our baseline findings are robust to industry-level price variation and macroeconomic shocks. While the coefficient estimates remain directionally consistent, their statistical significance weakens under industry-year fixed effects. This attenuation is likely driven by increased measurement error, as controlling for industry-year fixed effects absorbs a portion of the cross-industry variation in robot exposure—one of the primary sources of identifying variation in our design.

(ii) Pre-trend tests. To address concerns regarding pre-existing trends, we verify that our measure of robot exposure in European countries is not systematically related to past trends in innovation activities between 1988 and 2003, a period preceding the rapid advancements in robotics technology. We employ two approaches: First, as a placebo test, we examine whether robot exposure during the 2004–2019 period influences firms' innovation shift and productivity between 1988 and 2003. The regression results, presented in Panel A of Table A13, indicate that the estimated coefficients for innovation shift are not statistically significant. Second, we directly control for changes in firms' innovation activities between 1988 and 2003 by including them as additional explanatory variables in our baseline specifications. Panel B shows that this control has little effect on the baseline estimates.

5.1.2 Alternative IV constructions

We employ two alternative methods to construct the instrumental variable.

(i) Incorporating additional EU countries. The first method expands the robot usage data to include additional EU countries, as the baseline set of five countries may not fully capture broader automation trends across Europe. We modify the construction of the instrumental variable in two ways, following the approach of Acemoglu and Restrepo (2020).[17] First, we incorporate Germany

[17] As noted by Acemoglu and Restrepo (2020), robot adoption rates in the five baseline countries closely track those of the United States, while Germany's rate is significantly higher. Norway, Spain, and the United Kingdom exhibit adoption rates comparable to those in the United States.

into the five-country average. Second, we expand the average to include nine European countries with advanced automation: Finland, Sweden, Denmark, Italy, France, Germany, Norway, Spain, and the United Kingdom.[18] These countries are selected based on the availability of comprehensive industry statistics in the IFR database and collectively account for 41% of the global industrial robot market.

(ii) Adjusting for time-related confounds. In the second approach, we address potential time-related trends or unobserved temporal factors that may affect the instrumental variable. Specifically, we regress the baseline IV on time-fixed effects and use the resulting residuals as an adjusted measure to re-estimate the main results.

The results, presented in Table A14, confirm the robustness of our findings; the observed effects are not driven by the choice of European countries or by spurious time-related variation.

5.1.3 Construction of the shift-share IV

To mitigate the potential issue that both robot exposure in the U.S. firms and the number of robots per unit of labor across EU countries may vary over time due to common underlying factors, we adopt a shift-share instrumental variable (SSIV) approach in this section. Specifically, we adopt the SSIV framework of Borusyak, Hull, and Jaravel (2022), rather than that of Goldsmith-Pinkham, Sorkin, and Swift (2020), and treat robot usage per worker across EU countries as a quasi-experimental shock. This strategy, which relies on exogenous shocks rather than exogenous shares, is further motivated by Acemoglu and Restrepo (2022), who show that historical differences in industry-specific aging patterns generated heterogeneous robot exposure across industries. As a result, robot usage in EU countries is primarily driven by industry-specific characteristics that remain stable over time and are exogenous to individual firms' decisions. The extent of robot exposure for each firm is then determined by the annual sales distribution of firms across industries, integrating fixed industry characteristics with temporal variation in firms' sensitivity to industry-level robot adoption.

Constructing SSIV. Specifically, we modify the baseline IV by replacing its time-varying component—representing the annual proportion of robot usage to employment across EU countries—with

[18] Since EU KLEMS does not provide data for Norway, we impute its employment distribution using data from Denmark, Finland, and Sweden.

a fixed average proportion calculated over the period from 2000 to 2004, the five years preceding the start of the sample. The construction of this IV is as follows:

$$Robot\ Exposure\ SSIV_{it} = \sum_{k} \frac{Sales_{ikt}}{Sales_{it}} \times \frac{Robots^{EU}_{k,\text{00-04}}}{Workers^{EU}_{k,\text{00-04}}}, \tag{14}$$

where $Robots^{EU}_{k,\text{00-04}}$ and $Workers^{EU}_{k,\text{00-04}}$ represent the five-year average counts of industrial robots and workers, respectively, in industry k during the period 2000–2004 across EU countries, including Denmark, France, Finland, Italy, and Sweden.

Validity tests. We conduct a series of validity tests to ensure that the exogenous shocks satisfy the requirements of quasi-randomness and large-sample consistency. To assess sample adequacy, we calculate the effective sample size across industries and time periods, obtaining a value of 63.38, as determined by the inverse of the Herfindahl-Hirschman Index. This result indicates substantial industry-level variation.[19] Additionally, we evaluate the within-group correlation of shocks and confirm that the condition of sufficiently uncorrelated shocks is met.

Falsification tests. We conduct falsification tests by regressing potential industry-level and firm-level confounders on the instrument while employing exposure-robust inference to account for inherent dependencies in the data.[20] Following Acemoglu and Restrepo (2020), we control for several industry-level factors, including the share of high-tech equipment in total investment (*High-tech/Invest*), the ratio of input to value added (*Input/Value-added*), the log of real wages (*Ln(Wage)*), wage per worker (*Wage/Worker*), the log of employment (*Ln(Worker)*), and the ratio of workers to value added (*Worker/Value-added*). At the firm level, we include the baseline controls and a dummy indicating the adoption of an ES system. We further conduct a firm-level pre-trend test by regressing pre-2004 variables (1988–2003) on the shift–share instrument. If shocks are randomly assigned across industries, they should not predict these predetermined variables. Panel A of Table A15 confirms that none of the coefficients are statistically significant.

[19] According to Borusyak, Hull, and Jaravel (2022), even an effective sample size as small as 20 can yield reasonably accurate statistical inference in industry-level regressions.

[20] Following Borusyak, Hull, and Jaravel (2022), we estimate the coefficients using equivalent industry-level regressions to obtain valid exposure-robust standard errors.

Results. Panel B of Table A15 reports the SSIV estimates. Following Borusyak, Hull, and Jaravel (2022), we estimate the coefficients using equivalent industry-level regressions to obtain valid exposure-robust standard errors. Consistent with the methodology described above, standard errors are clustered at the industry level to account for potential within-industry correlation. The results are generally consistent with our main findings, confirming that increased robot exposure leads to a shift in innovation towards AI and greater R&D investment.

5.2 Robustness of Text-Based Robot Exposure Measures

5.2.1 Alternative measures of robot exposure

We explore several alternative approaches to measuring robot exposure.

(i) Refining and reweighting keywords. We refine keywords using a two-stage regression process, following Cong et al. (2024). The first stage employs a LASSO regression for variable selection: $y_{kt} = \sum_{w \in W} \beta_w x_{wkt} + \varepsilon_{kt}$, where y_{kt} denotes robot exposure in industry k at time t, measured as the number of industrial robots installed per thousand workers. The variable x_{wkt} represents the intensity of keyword w in industry k and year t, and $w \in W$ is an element of the predefined keyword set W, as presented in Table A2.

In the second stage, we retain keywords with non-zero coefficients from the LASSO regression and re-estimate their weights using OLS. Let W^* denote the refined keyword set and β_w the corresponding OLS coefficient for keyword $w \in W^*$. Each firm's robot exposure in year t is then computed as $\textit{Robot Exposure } A_{it} = \sum_{w \in W^*} \beta_w x_{wit}$, where x_{wit} represents the scaled occurrence of keyword w in firm i's 10-K filing for year t.

(ii) Adjusting for time-related confounds. A potential concern is that our results may reflect common time trends or co-integration rather than a causal effect of robot exposure on innovation. To address this issue, we regress robot exposure on time-fixed effects and use the residuals as the main explanatory variable to remove time-specific patterns and control for temporal confounding.

(iii) Adjusting for industry-related confounds. Our findings could also be driven by systematic differences across industries rather than by firm-level heterogeneity. To mitigate this risk, we regress robot exposure on industry-fixed effects and use the residuals as the main explanatory variable,

thereby eliminating industry-level variation and isolating within-industry differences across firms.

(iv) Applying rolling weighted average. Since robots are capital assets that depreciate over time, relying on robot-related mentions from a single year may not accurately reflect their actual value and utilization. To address this concern, we construct a rolling five-year weighted average of robot exposure. Specifically, we assign weights of 0.2, 0.4, 0.6, 0.8, and 1 to the firm's robot exposure from year $t-4$ to t, allowing previous investments to influence the exposure measure with gradually declining impact.

(v) Applying the ratio of robots to employment. Finally, we measure robot exposure as the ratio of robots to employment, following the methodology of Acemoglu and Restrepo (2020). Firm-level robot exposure is then defined as: $\textit{Robot Exposure } D_{it} = \sum_k \frac{Sales_{ikt}}{Sales_{it}} \times \frac{Robots_{kt}}{Workers_{kt}}$, where $Sale_{ikt}$ represents the annual sales of firm i in segment (industry) k at year t, and $Sale_{it}$ denotes the total sales of firm i in year t. The term $\frac{Robots_{kt}}{Workers_{kt}}$ captures industry-level robot exposure for industry k at year t.

5.2.2 Controlling for 10-K text readability

Because our measure of robot exposure is constructed from firms' 10-K filings, the results may be influenced by variations in textual complexity or disclosure style. Following F. Li (2008), we calculate the readability of each firm-year 10-K using the Flesch Reading Ease (FRE) score,[21] which is then included as an additional control variable in our regressions.

Table A16 presents the results, indicating that the magnitude and statistical significance of the core variable are comparable to the baseline estimates, lending support to our findings.

5.3 Effects of the Supreme Court's *Alice* Decision

We examine the impact of the Supreme Court's 2014 decision in *Alice Corp. v. CLS Bank International* on the patentability of AI-related inventions. The ruling significantly reshaped U.S. patent law by holding that inventions based on "abstract ideas" are not patentable unless they contain an "inventive concept." As a result, AI innovations involving abstract algorithms or processes may face

[21] The score ranges from 0 to 100, with higher values indicating greater readability. It is calculated as $\text{FRE} = 206.835 - 1.015 \times (\text{words}/\text{sentences}) - 84.6 \times (\text{syllables}/\text{words})$. While standard readability metrics may overstate complexity in financial disclosures, the FRE remains a widely used benchmark facilitating comparability with prior literature.

greater challenges in obtaining patent protection in the post-*Alice* era.

Effects on AI patent approvals. We first examine whether the *Alice* decision increased the difficulty of obtaining AI patents. Using USPTO pre-grant data, we identify AI-related applications through abstract analysis and merge them with granted patents covering 2004–2019. We measure the impact using two variables: $Dummy(Grant)$, an indicator for whether a patent is granted, and $Ln(Grant\ Lag)$, the log of the days from application to approval (plus one). To test differential effects, we include an interaction term, $Dummy(AI\ Patent) \times Dummy(After\ 2014)$, where $Dummy(AI\ Patent)$ equals 1 for AI patents and $Dummy(After\ 2014)$ equals 1 for patents filed in or after 2014.

Panel A of Table A17 presents the summary of results. Columns (1) and (2) show that the coefficient of the interaction term is negative but not statistically significant, indicating that there is no meaningful impact of *Alice* on AI patent approval rates. Column (3) examines the approval times for granted patents, finding that the interaction term is positive and significant at the 5% level, while $Dummy(AI\ Patent)$ is also positive. This suggests that AI patents take longer to approve than non-AI patents, with extended delays post-*Alice*. Column (4) provides robustness checks using a Poisson model, yielding consistent findings. These results indicate that while the *Alice* decision does not significantly affect approval rates, it does lengthen approval times for AI patents.

Robustness checks: The *Alice* decision. We conduct robustness checks examining whether the *Alice* decision influences the AI shift. Panel B of Table A17 shows that the ruling does not significantly affect the results for AI-directed innovation or innovation quality in the patent–firm–year regressions.

5.4 Alternative Approaches to Sample Constructions

For robustness, we consider three alternative sample constructions:

(i) Excluding the impacts of other Internet technologies. Given that emerging technologies such as big data and cloud computing are closely associated with AI, they may also confound our main regression results. To mitigate the influence of such technological shocks, we adopt two exclusion strategies. First, we exclude firms operating in the Internet industry, as they are more likely

to be directly affected by developments in cloud computing and other internet-based technologies.[22] Second, we exclude firms that adopted Enterprise Systems (ES), a broad category of IT investments known for their advanced capabilities and wide functional scope (Dorantes et al. 2013). We identify ES adopters using ES-related keywords in the LexisNexis database.

(ii) Excluding firms with significant overseas operations. Many U.S. firms in our sample maintain overseas production units, which may influence our main findings. To address this concern, we exclude firms whose overseas sales exceed 50% of their total sales.

(iii) Excluding firms with extremely high robot exposure. To ensure that our results are not driven by extreme automation levels in a small subset of firms or industries, we exclude the top 1% of firms in terms of the automation index within each industry-year.

The results of the tests based on these adjusted samples are reported in Table A18. The estimates show a similar pattern to our baseline findings.

6 A Simple Equilibrium Model of Innovation Shift

We now develop a dynamic equilibrium model to rationalize the observed innovation shift in response to rising robot exposure. Appendix 7 contains the proofs for all propositions.

Model Setup

A continuous-time economy with $t \in [0,\infty)$ features a continuum of firms indexed by $f \in [0,1]$. Following Acemoglu and Restrepo (2019), we model the degree of automation within firms as follows: Each firm operates a continuum of tasks indexed by $s \in [0,1]$. A given firm f automates tasks in the range $s \in [0,\theta_f]$, while tasks $s \in [\theta_f,1]$ remain non-automated. Labor is supplied inelastically, with one unit available in the market.

Firms undertake two types of innovations: one aimed at enhancing the productivity of automated tasks and the other at improving the productivity of non-automated tasks. For tasks in $[0,\theta_f]$, firms innovate, patent their innovations, and leverage these patented technologies to enhance task productiv-

[22] Specifically, we define internet firms as those with four-digit SIC codes between 7371 and 7379, which correspond to computer programming, software publishing, data processing, and other computer-related services.

ity. Innovations related to automated tasks are integrated with machines to produce the final product:

$$\frac{1}{1-\beta} q_M(f)^{\beta} k_M^{1-\beta}(s \mid f), \tag{15}$$

where $q_M(f)$ represents the technology used in automated tasks, and $k_M(s \mid f)$ denotes the units of machines employed in task s by firm f. The production of one unit of a machine requires $\frac{1}{\chi_M}$ units of the final good. For non-automated tasks, firm f combines innovations applicable to these tasks with labor to produce its final products according to

$$\frac{1}{1-\beta} q_L^{\beta}(f) k_L^{1-\beta}(s \mid f), \tag{16}$$

where $q_L(f)$ denotes the technology used in non-automated tasks, and $k_L(s|f)$ is the labor-driven production input. Each unit of $k_L(s|f)$ can be produced at a cost of $\frac{1}{\chi_L \bar{q}_L}$, with $\bar{q}_L = \int_0^1 (1-\theta_f) q_L(f) df$. We normalize the production of $k_L(s|f)$ by $\bar{q}_L$ to ensure that the stationary equilibrium remains linear. We normalize the price of the final product to one for each time $t \in [0,\infty)$. Firm f chooses $k_M(s|f)$ and $k_L(s|f)$ to maximize instantaneous profit conditional on the latest technology,

$$\begin{aligned} \Pi_f(q_M, q_L) = \max_{k_M(s|f), k_L(s|f)} & \frac{1}{1-\beta} \int_0^{\theta_f} q_M^{\beta}(f) k_M^{1-\beta}(s \mid f)\, ds - \frac{1}{\chi_M} \int_0^{\theta_f} k_M(s \mid f) ds \\ & + \frac{1}{1-\beta} \int_{\theta_f}^1 q_L^{\beta}(f) k_L^{1-\beta}(s \mid f) ds - \int_{\theta_f}^1 \frac{w k_L(s \mid f)}{\bar{q}_L \chi_L} ds, \end{aligned} \tag{17}$$

where w represents the wage for labor, determined by labor market clearing.

Labor market clearing implies that $w = \bar{q}_L \chi_L^{1-\beta}$. The firm's first-order conditions yield a profit function that is linear in $q_M(f)$ and $q_L(f)$:

$$\Pi_f = \alpha_M \theta_f q_M(f) + \alpha_L (1-\theta_f) q_L(f), \tag{18}$$

with $\alpha_M = \frac{\beta}{1-\beta} \chi_M^{\frac{1-\beta}{\beta}}$ and $\alpha_L = \frac{\beta}{1-\beta} \chi_L^{1-\beta}$.

The normalization by $\bar{q}_L$ ensures a linear stationary equilibrium. Solving the task-level FOCs delivers a firm profit function linear in (q_M, q_L), which is what makes the linear value-function guess (below) valid.

The marginal instantaneous return of innovation in automated tasks is given by $\frac{\partial \Pi_f}{\partial q_M(f)} = \alpha_M \theta_f := \gamma_f$, while the marginal instantaneous profit for non-automated tasks is $\alpha_L(1-\theta_f)$. The parameter θ_f captures the *extensive margin*: the larger the θ_f, the greater the fraction of tasks that are automated. Consequently, innovation in automated processes benefits more tasks. The term α_M represents the *intensive margin*: a larger α_M, caused by a larger χ_M, implies a lower marginal cost of using robots per task, k_M, leading to increased robot usage in each automated production line, thereby enhancing the marginal return of innovation in automated tasks.

Together, γ_f captures the total marginal return to innovation in automated tasks, reflecting both extensive and intensive margins. Empirically, we interpret γ_f as the automation exposure of firm f, which depends on these same margins. In our theoretical framework, we assume that α_M is identical across firms, while firms differ in their degree of automation, θ_f. In practice, however, various factors can induce variation in α_M. For instance, an increase in the cost of materials that complement robot use on production lines can be interpreted as a decline in α_M. Similarly, firms located in different regions may face heterogeneous wage pressures; regions with high wages can be interpreted as experiencing a high α_M.

6.1 Innovation

We follow Akcigit and Kerr (2018) in modeling innovation as a stochastic arrival process, but extend their framework to accommodate heterogeneous innovation across task types: automated and non-automated tasks. Let $k \in \{M, L\}$ denote the task type, where M refers to automated tasks and L to non-automated tasks. Firms can invest in R&D to enhance the arrival rate z_k of innovations that improve the technology level q_k used in task k.

The cost of achieving an arrival rate z_k for technology q_k is given by

$$C_k(z_k, q_k) = C_k(z_k)\, q_k,$$

where the cost function satisfies standard properties: $C_k(0) = 0$, $C_k'(0) = 0$, $C_k''(z_k) > 0$, and there exists a finite upper bound $\bar{z}$ such that $C_k(\bar{z}) = \infty$. Empirically, data generated through automated tasks can be interpreted as a channel that effectively reduces the cost of future innovation, shifting the

$C_M(\cdot)$ downward. By providing structured information on how robots are used in production, firms gain insights that facilitate subsequent technological improvements, lower experimentation costs, and accelerate algorithmic learning in automated processes.

To reflect the time-varying nature of technology, we denote the technology level of firm f in task k at time t by $q_k(t \mid f)$, which may be interchangeably written as $q_k(f)$ when time is not the focus. Given arrival rate z_k, the technology level evolves as:

$$q_k(t+\Delta t \mid f) = \begin{cases} (1+\lambda_k) \cdot q_k(t \mid f) & \text{with probability } z_k \Delta t, \\ q_k(t \mid f) & \text{with probability } 1 - z_k \Delta t, \end{cases} \tag{19}$$

where λ_k denotes the step size of technological improvement from a single innovation.[23]

6.2 R&D Decision

For a firm f, denoting its value as V, we consider the stationary solution to the economy and the Markov Perfect Equilibrium, where the value function depends on the state variables q_M and q_L. We have a Bellman equation:

$$\begin{aligned} rV = \max_{z_M, z_L} \{ & \Pi_f - C_M(z_M) q_M - C_L(z_L) q_L + z_M [V(q_M(1+\lambda_M), q_L) - V(q_M, q_L)] \\ & + z_L [V(q_M, q_L(1+\lambda_L)) - V(q_M, q_L)] \} \end{aligned} . \tag{20}$$

Because instantaneous profits are linear in (q_M, q_L) and R&D costs scale with (q_M, q_L), the HJB admits a linear stationary solution. This pins down closed-form FOCs and comparative statics that generate testable sign predictions.

Here, we omit the subscript f for notational simplicity without loss of clarity. The equation is intuitive: on the right-hand side, firm f receives the instantaneous cash flow, which consists of the instantaneous profit Π_f from production minus the R&D expenditure, along with the expected future cash flow from innovation. The terms $z_M [V(q_M(1+\lambda_M), q_L) - V(q_M, q_L)]$ and $z_L [V(q_M, q_L(1+\lambda_L)) - V(q_M, q_L)]$

[23] That is, the firm experiences no change in q_k with probability $1 - z_k \Delta t$, and a discrete jump of size λ_k with probability $z_k \Delta t$. The probability of multiple innovations occurring within Δt is of order $o(\Delta t)$, and is therefore negligible in the continuous-time limit.

represent the net expected value from innovation in automated and non-automated tasks, respectively. The left-hand side of the equation accounts for the firm's value discounted at rate r.

Proposition 6.1. *The unique solution to Equation (20) is given by* $V(q_M, q_L) = A_M q_M + A_L q_L$, *where* A_M *and* A_L *are determined by:*

$$\begin{cases} rA_M = \alpha_M \theta_f - C_M(z_M^*) + z_M^* A_M \lambda_M, \\ rA_L = (1-\theta_f)\alpha_L - C_L(z_L^*) + z_L^* A_L \lambda_L. \end{cases} \tag{21}$$

The arrival rates z_M^* *and* z_L^* *chosen by firm* f *are determined from the first-order conditions:*

$$C_M'(z_M^*) = A_M \lambda_M, \quad \text{and} \quad C_L'(z_L^*) = \lambda_L A_L. \tag{22}$$

6.3 Automation and Innovation Shift

We examine how automation affects innovation in automated and non-automated tasks. Let overall automation be defined as $\gamma_f = \theta_f \times \alpha_M$, where θ_f captures the extensive margin (task coverage) and α_M the intensive margin (degree of automation). We analyze the comparative statics of innovation arrival, z_k^*, with respect to γ for $k \in \{M, L\}$.

Combine (i) the HJB-implied linear value system for (A_M, A_L) and (ii) the FOCs for (z_M, z_L). A convenient approach is to substitute $A_M = C_M'(z_M^*)/\lambda_M$ into the linear equation for A_M, yielding an implicit equation in z_M^* and γ_f, then apply the implicit function theorem. This yields $\partial z_M^*/\partial \gamma_f > 0$ and $\partial z_L^*/\partial \gamma_f < 0$ under standard regularity conditions (including finiteness of firm value, which requires $r > z_k^* \lambda_k$). For expositional simplicity, we present the sign result and a compact expression below.

Proposition 6.2. *An increase in automation raises innovation in automated tasks but reduces innovation in non-automated tasks:*

$$\frac{\partial z_M^*}{\partial \gamma_f} > 0, \quad \frac{\partial z_L^*}{\partial \gamma_f} < 0. \tag{23}$$

A more explicit form is

$$\frac{\partial z_M^*}{\partial \gamma} = \frac{\lambda_M}{(r - \lambda_M z_M^*) C_M''(z_M^*)} > 0, \quad \frac{\partial z_L^*}{\partial \gamma} = -\frac{\lambda_L}{(r - \lambda_L z_L^*) C_L''(z_L^*)} \cdot \frac{\alpha_L}{\alpha_M} < 0, \tag{24}$$

which reduces to the commonly used approximation with denominator r when $\lambda_k z_k^* < r$.

This proposition shows that an increase in automation exposure encourages firms to allocate more effort toward innovation in automated tasks (M), while reducing innovation in non-automated tasks (L), leading to a rise (decline) in the arrival rate in automated (non-automated) tasks. It formalizes the intuitive trade-off: as automation raises returns to innovation in automated activities, resources are reallocated away from non-automated areas.

Interpretation of γ. We empirically proxy automation γ with firm-level robot exposure. The two components—θ_f (the scope of automation) and α_M (automation intensity)—map naturally to this measure. Innovation in automated tasks (z_M) is proxied using AI-related patents, while innovation in non-automated tasks (z_L) is proxied using other patents.

Role of cost curvature $C_k''(\cdot)$. Assume a quadratic cost function $C_k(z) = \frac{c_k}{2} z^2$. Under this specification, firms with lower c_M exhibit greater responsiveness to increases in γ, reallocating innovation more aggressively toward AI. If prior AI engagement lowers c_M, then AI-experienced firms will disproportionately shift R&D toward AI in response to automation.

Role of innovation step size λ_M. The parameter λ_M captures the productivity impact of a single AI innovation. Firms with higher λ_M benefit more from each AI breakthrough and hence reallocate R&D toward AI more strongly when automation rises.

Role of data availability. An increase in data availability enhances the productivity of AI algorithms, effectively lowering c_M. Thus, more data facilitates innovation in automated tasks, further amplifying the shift toward AI-intensive innovation.

Total R&D, arrival rates, and the innovation J-curve To connect the model to the empirically observed innovation "J-curve," we focus on total R&D expenditure and the evolution of arrival rates (z_M, z_L) across automated and non-automated tasks.

Consider a rise in AI-oriented innovation z_M^* and a corresponding decline in non-automated innovation z_L^*. The total R&D expenditure changes according to the full derivative:

$$d\text{Expenditure} = C'_M(z_M^*)\, dz_M + C'_L(z_L^*)\, dz_L. \tag{25}$$

Under standard convexity assumptions on $C_k(\cdot)$, there exist conditions where the total arrival rate declines:

$$dz_M + dz_L < 0, \tag{26}$$

while total R&D expenditure rises:

$$d\text{Expenditure} > 0, \tag{27}$$

because the marginal cost of R&D in automated tasks increases faster than the reduction in non-automated R&D.

Over time, as data accumulates and the marginal cost of automated-task R&D (C_M) declines—through learning, algorithmic improvement, or experience—the effective arrival rate of automated-task innovation increases. This generates the characteristic **J-curve** in total innovation: a temporary slowdown in overall arrivals followed by a rebound as the accumulated knowledge reduces costs and accelerates innovation.

Finally, we need to remind the limits of our simple model. This stylized model does not: (i) model the high-dimensional technology space underlying semantic embeddings; (ii) microfound the citation-network distance/closeness measures; (iii) endogenize the evolution of θ_f (automation adoption) within the model; (iv) explicitly model multi-stage research vs development beyond the reduced-form (π_k, τ_k) wedge; or (v) deliver quantitative magnitudes without calibration. Its objective is to provide transparent sign predictions and heterogeneity predictions that map to the empirical results.

7 Conclusion

In recent year, corporate innovations appear to have slowed down despite increases in R&D expenditure. We find evidence that this is a manifestation of firms' endogenous responses to rising automation and robot exposure around the globe. We argue that automation and AI innovations are complements

because the former generates data for the latter, which in turn improves the former's efficiency. As firms shift corporate innovations towards AI-oriented ones that entail greater generality and originality, and thus greater cost, they naturally witness significant immediate declines in patent output but rises in innovation inputs, with potential long-run benefits. Such automation-induced innovation shifts are more significant for firms with more AI-related research experience or generate more data, and can be rationalized in a simple dynamic model.

Our study not only provides an explanation for the puzzling divergence between R&D spending and patent output, but also constitutes an initial investigation into the broader question of how one innovation or technology adoption induces other innovations in the short run and over the long term, which warrants further research. Our findings also inform business strategies and regulatory policies. In particular, policy makers and firms may need to be more patient with investments associated with innovation shifts.

References

Abis, S. and L. Veldkamp (2024). "The changing economics of knowledge production", *The Review of Financial Studies*, 37, 89–118.

Acemoglu, D., U. Akcigit, and W. R. Kerr (2016). "Innovation network", *Proceedings of the National Academy of Sciences*, 113, 11483–11488.

Acemoglu, D. and J. Linn (2004). "Market size in innovation: Theory and evidence from the pharmaceutical industry", *The Quarterly Journal of Economics*, 119, 1049–1090.

Acemoglu, D. and P. Restrepo (2019). "Automation and new tasks: How technology displaces and reinstates labor", *Journal of Economic Perspectives*, 33, 3–30.

– (2020). "Robots and jobs: Evidence from US labor markets", *Journal of Political Economy*, 128, 2188–2244.

– (2022). "Demographics and automation", *The Review of Economic Studies*, 89, 1–44.

Acharya, V. V., R. P. Baghai, and K. V. Subramanian (2014). "Wrongful discharge laws and innovation", *The Review of Financial Studies*, 27, 301–346.

Acharya, V. V. and K. V. Subramanian (2009). "Bankruptcy codes and innovation", *The Review of Financial Studies*, 22, 4949–4988.

Aghion, P. and P. Howitt (1992). "A model of growth through creative destruction", *Econometrica*, 60, 323.

Akcigit, U., J. Grigsby, T. Nicholas, and S. Stantcheva (2022). "Taxation and innovation in the twentieth century", *The Quarterly Journal of Economics*, 137, 329–385.

Akcigit, U. and W. R. Kerr (2018). "Growth through heterogeneous innovations", *Journal of Political Economy*, 126, 1374–1443.

Arora, A., S. Belenzon, and A. Patacconi (2018). "The decline of science in corporate R&D", *Strategic Management Journal*, 39, 3–32.

Autor, D. and A. Salomons (2018). "Is automation labor-displacing? Productivity growth, employment, and the labor share", *NBER Working Paper No. w24871*.

Babina, T., A. Fedyk, A. He, and J. Hodson (2024). "Artificial intelligence, firm growth, and product innovation", *Journal of Financial Economics*, 151, 103745.

Bajari, P., V. Chernozhukov, A. Hortaçsu, and J. Suzuki (2019). "The impact of big data on firm performance: An empirical investigation", *AEA Papers and Proceedings*, 109, 33–37.

Bekamiri, H., D. S. Hain, and R. Jurowetzki (2021). "PatentSBERTa: A deep NLP based hybrid model for patent distance and classification using augmented SBERT", *arXiv preprint arXiv:2103.11933*.

Bena, J. and K. Li (2014). "Corporate innovations and mergers and acquisitions", *The Journal of Finance*, 69, 1923–1960.

Beraja, M., D. Y. Yang, and N. Yuchtman (2023). "Data-intensive innovation and the state: Evidence from AI firms in China", *The Review of Economic Studies*, 90, 1701–1723.

Berg, A., E. F. Buffie, and L.-F. Zanna (2018). "Should we fear the robot revolution?(The correct answer is yes)", *Journal of Monetary Economics*, 97, 117–148.

Bloom, N., S. J. Davis, and Y. Zhestkova (2021). "Covid-19 shifted patent applications toward technologies that support working from home", *AEA Papers and Proceedings*, 111, 263–66.

Bloom, N., C. I. Jones, J. Van Reenen, and M. Webb (2020). "Are ideas getting harder to find?", *American Economic Review*, 110, 1104–1144.

Bloom, N., M. Schankerman, and J. Van Reenen (2013). "Identifying technology spillovers and product market rivalry", *Econometrica*, 81, 1347–1393.

Borusyak, K., P. Hull, and X. Jaravel (2022). "Quasi-experimental shift-share research designs", *The Review of Economic Studies*, 89, 181–213.

Brynjolfsson, E. and K. McElheran (2016). "The rapid adoption of data-driven decision-making", *American Economic Review*, 106, 133–139.

Brynjolfsson, E., D. Rock, and C. Syverson (2021). "The productivity J-curve: How intangibles complement general purpose technologies", *American Economic Journal: Macroeconomics*, 13, 333–372.

Caselli, F. and A. Manning (2019). "Robot arithmetic: New technology and wages", *American Economic Review: Insights*, 1, 1–12.

Cheng, X., E. Lyandres, K. Zhou, and T. Zhou (2024). "Labor-replacing automation and finance", *Management Science*.

Chuang, H.-C., P.-H. Hsu, C.-M. Kuan, and J.-C. Yang (2024). "Limitation of firm fixed effects models and the missing R&D-patent relation: New methods and evidence", *SSRN Working Paper*, 4636846.

Cockburn, I. M., R. Henderson, and S. Stern (2018). "The impact of artificial intelligence on innovation: An exploratory analysis", *The Economics of Artificial Intelligence: An Agenda*.

Cohn, J. B., Z. Liu, and M. I. Wardlaw (2022). "Count (and count-like) data in finance", *Journal of Financial Economics*, 146, 529–551.

Cong, L. W., T. Liang, X. Zhang, and W. Zhu (2024). "Textual factors: A scalable, interpretable, and data-driven approach to analyzing unstructured information", *NBER Working Paper No. w33168*.

Costinot, A., D. Donaldson, M. Kyle, and H. Williams (2019). "The more we die, the more we sell? A simple test of the home-market effect", *The Quarterly Journal of Economics*, 134, 843–894.

Cui, Y., Y. Lu, and W. Zhu (2025). "Innovation network and mergers and acquisitions", *Working Paper*.

De Bettignies, J.-E., B. Gainullin, H. F. Liu, and D. T. Robinson (2018). "The effects of downstream competition on upstream innovation and licensing", *NBER Working Paper No. w25166*.

Dixon, J., B. Hong, and L. Wu (2021). "The robot revolution: Managerial and employment consequences for firms", *Management Science*, 67, 5586–5605.

Dorantes, C.-A., C. Li, G. F. Peters, and V. J. Richardson (2013). "The effect of enterprise systems implementation on the firm information environment", *Contemporary Accounting Research*, 30, 1427–1461.

Einiö, E., J. Feng, and X. Jaravel (2019). "Social push and the direction of innovation", *SSRN Working Paper*, 3383703.

Eisfeldt, A. L., G. Schubert, and M. B. Zhang (2023). "Generative AI and firm values", *NBER Working Paper No. w31222*.

Fang, H., X. Gu, H. Yan, and W. Zhu (2025). "AI Inventions and AI Talent Migrations: US and China".

Fang, L. H., J. Lerner, and C. Wu (2017). "Intellectual property rights protection, ownership, and innovation: Evidence from China", *The Review of Financial Studies*, 30, 2446–2477.

Finkelstein, A. (2004). "Static and dynamic effects of health policy: Evidence from the vaccine industry", *The Quarterly Journal of Economics*, 119, 527–564.

Frésard, L., G. Hoberg, and G. M. Phillips (2020). "Innovation activities and integration through vertical acquisitions", *The Review of Financial Studies*, 33, 2937–2976.

Giczy, A. V., N. A. Pairolero, and A. A. Toole (2022). "Identifying artificial intelligence (AI) invention: A novel AI patent dataset", *The Journal of Technology Transfer*, 47, 476–505.

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift (2020). "Bartik instruments: What, when, why, and how", *American Economic Review*, 110, 2586–2624.

Gourieroux, C., A. Monfort, and A. Trognon (1984). "Pseudo maximum likelihood methods: Theory", *Econometrica: Journal of the Econometric Society*, 681–700.

Graetz, G. and G. Michaels (2018). "Robots at work", *Review of Economics and Statistics*, 100, 753–768.

Griffith, R., S. Redding, and J. Van Reenen (2003). "R&D and absorptive capacity: Theory and empirical evidence", *Scandinavian Journal of Economics*, 105, 99–118.

Griliches, Z. (2009). "Productivity, R&D, and the data constraint", *The Economic Impact of Knowledge*.

Guerreiro, J., S. Rebelo, and P. Teles (2022). "Should robots be taxed?", *The Review of Economic Studies*, 89, 279–311.

Hall, B. H., A. B. Jaffe, and M. Trajtenberg (2001). "The NBER patent citation data file: Lessons, insights and methodological tools", *NBER Working Paper No. w8498*.

Hill, R. et al. (2025). "The pivot penalty in research", *Nature*, 1–8.

Hirshleifer, D., P.-H. Hsu, and D. Li (2018). "Innovative originality, profitability, and stock returns", *The Review of Financial Studies*, 31, 2553–2605.

Hoberg, G. and V. Maksimovic (2015). "Redefining financial constraints: A text-based analysis", *The Review of Financial Studies*, 28, 1312–1352.

Hopenhayn, H. and F. Squintani (2021). "On the direction of innovation", *Journal of Political Economy*, 129, 1991–2022.

Hsu, P.-H., X. Tian, and Y. Xu (2014). "Financial development and innovation: Cross-country evidence", *Journal of Financial Economics*, 112, 116–135.

Humlum, A. (2019). "Robot adoption and labor market dynamics", *Working Paper*.

Jones, B. F. (2009). "The burden of knowledge and the "death of the renaissance man": Is innovation getting harder?", *The Review of Economic Studies*, 76, 283–317.

Jumper, J. et al. (2021). "Highly accurate protein structure prediction with AlphaFold", *Nature*, 596, 583–589.

Kerr, W. R. (2008). "Ethnic scientific communities and international technology diffusion", *The Review of Economics and Statistics*, 90, 518–537.

Koch, M., I. Manuylov, and M. Smolka (2021). "Robots and firms", *The Economic Journal*, 131, 2553–2584.

Kogan, L., D. Papanikolaou, A. Seru, and N. Stoffman (2017). "Technological innovation, resource allocation, and growth", *The Quarterly Journal of Economics*, 132, 665–712.

Li, F. (2008). "Annual report readability, current earnings, and earnings persistence", *Journal of Accounting and economics*, 45, 221–247.

Malamud, S. and F. Zucchi (2019). "Liquidity, innovation, and endogenous growth", *Journal of Financial Economics*, 132, 519–541.

McElheran, K., M.-J. Yang, Z. Kroff, and E. Brynjolfsson (2025). "The rise of industrial AI in America: Microfoundations of the productivity J-curve(s)", *U.S. Census Bureau Working Paper No. 25-27*.

Mezzanotti, F. and T. Simcoe (2023). "Research and/or development? Financial frictions and innovation investment", *NBER Working Paper No. w31521*.

Miric, M., N. Jia, and K. G. Huang (2023). "Using supervised machine learning for large-scale classification in management research: The case for identifying artificial intelligence patents", *Strategic Management Journal*, 44, 491–519.

Mokyr, J. (2011). "The gifts of Athena: Historical origins of the knowledge economy", *Princeton University Press*.

Romer, P. M. (1986). "Increasing returns and long-run growth", *Journal of Political Economy*, 94, 1002–1037.

Rothaermel, F. T. and D. L. Deeds (2004). "Exploration and exploitation alliances in biotechnology: A system of new product development", *Strategic Management Journal*, 25, 201–221.

Silva, J. S. and S. Tenreyro (2006). "The log of gravity", *The Review of Economics and Statistics*, 641–658.

Sunder, J., S. V. Sunder, and J. Zhang (2017). "Pilot CEOs and corporate innovation", *Journal of Financial Economics*, 123, 209–224.

WIPO (2019). "WIPO technology trends 2019: Artificial intelligence", *Geneva: World Intellectual Property Organization*.

Wu, L., L. Hitt, and B. Lou (2020). "Data analytics, innovation, and firm productivity", *Management Science*, 66, 2017–2039.

Zeira, J. (1998). "Workers, machines, and economic growth", *The Quarterly Journal of Economics*, 113, 1091–1117.

Zhang, M. B. (2019). "Labor-technology substitution: Implications for asset pricing", *The Journal of Finance*, 74, 1793–1839.

Zhu, C. (2019). "Big data as a governance mechanism", *The Review of Financial Studies*, 32, 2021–2061.

Figures and Tables

Figure I: The Relationship between Robot Exposure and Robot Used per Worker in the U.S.

This figure presents the relationship between estimated robot exposure—aggregated to industry-level annual averages—and the log-transformed industry-level annual averages of robot used per worker in the U.S.

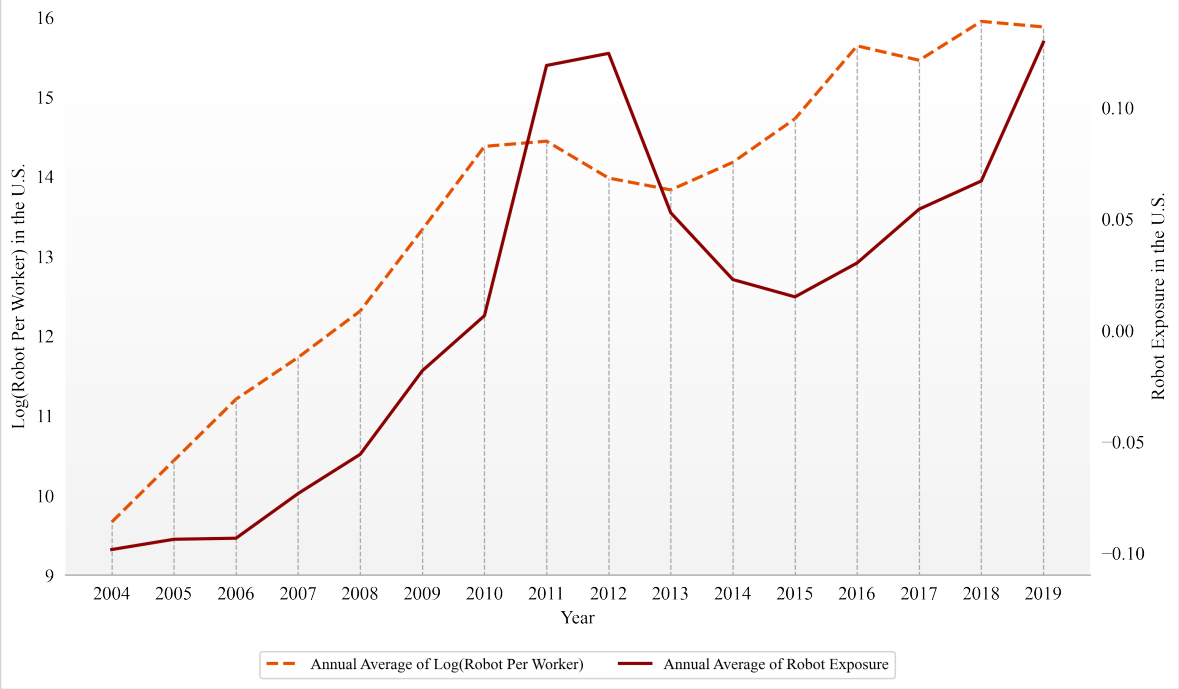
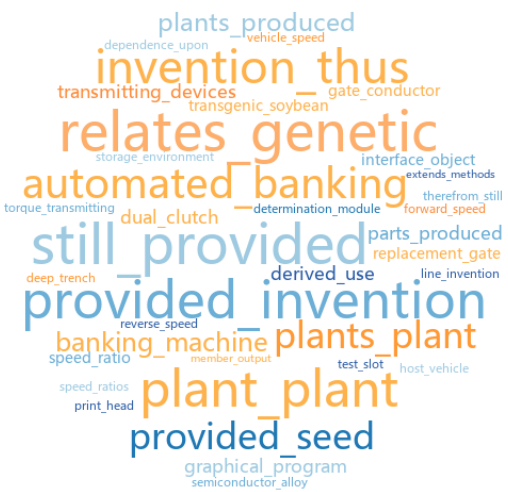


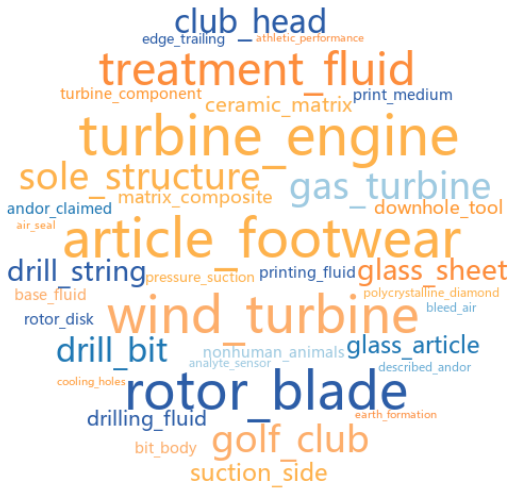
Figure II: Word Cloud of Relative Frequency Differences between High- and Low-Robot-Exposure Patents

This figure displays word clouds that visualize the relative frequency differences of n-grams in patent abstracts, comparing firms with high versus low levels of robot exposure. Patents are grouped according to whether the robot exposure of the corresponding firm is above or below the sample median. The analysis is conducted using samples from the years 2010, 2015, and 2019. For each group and year, n-gram frequencies are normalized relative to the baseline year of 2004. We then compute the ratio of normalized frequencies between the two groups and apply a natural log transformation to obtain group-level log-ratios of n-gram prominence. The resulting values are visualized using word clouds, where the font size of each n-gram reflects the magnitude of its relative frequency ratio—larger text indicates a greater relative frequency in one group compared to the other. Table A5 presents the top 20 n-grams with the highest relative frequency, corresponding to the figure.

(a) 2010



(b) 2015



(c) 2019

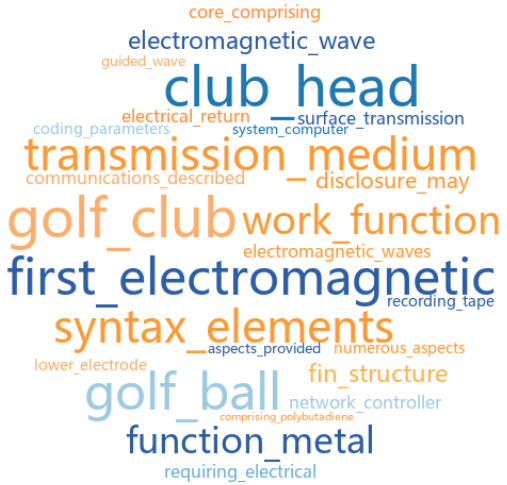


Figure III: Comparison of Relative Costs Between AI Patents and Non-AI Patents

This table presents a comparison of the relative costs of AI and non-AI patents across four dimensions: team size, labor input, inventor originality, and patent value, for patents filed between 2004 and 2019. *Team Size* is defined as the number of inventors contributing to a patent; *Labor Input* adjusts *Team Size* by addressing any potential double-counting of inventors; *Inventor Originality* is measured as one minus the sum of the shares of patents filed by the inventor in any given technological class; *Patent Value* follows the approach proposed by Kogan et al. (2017).

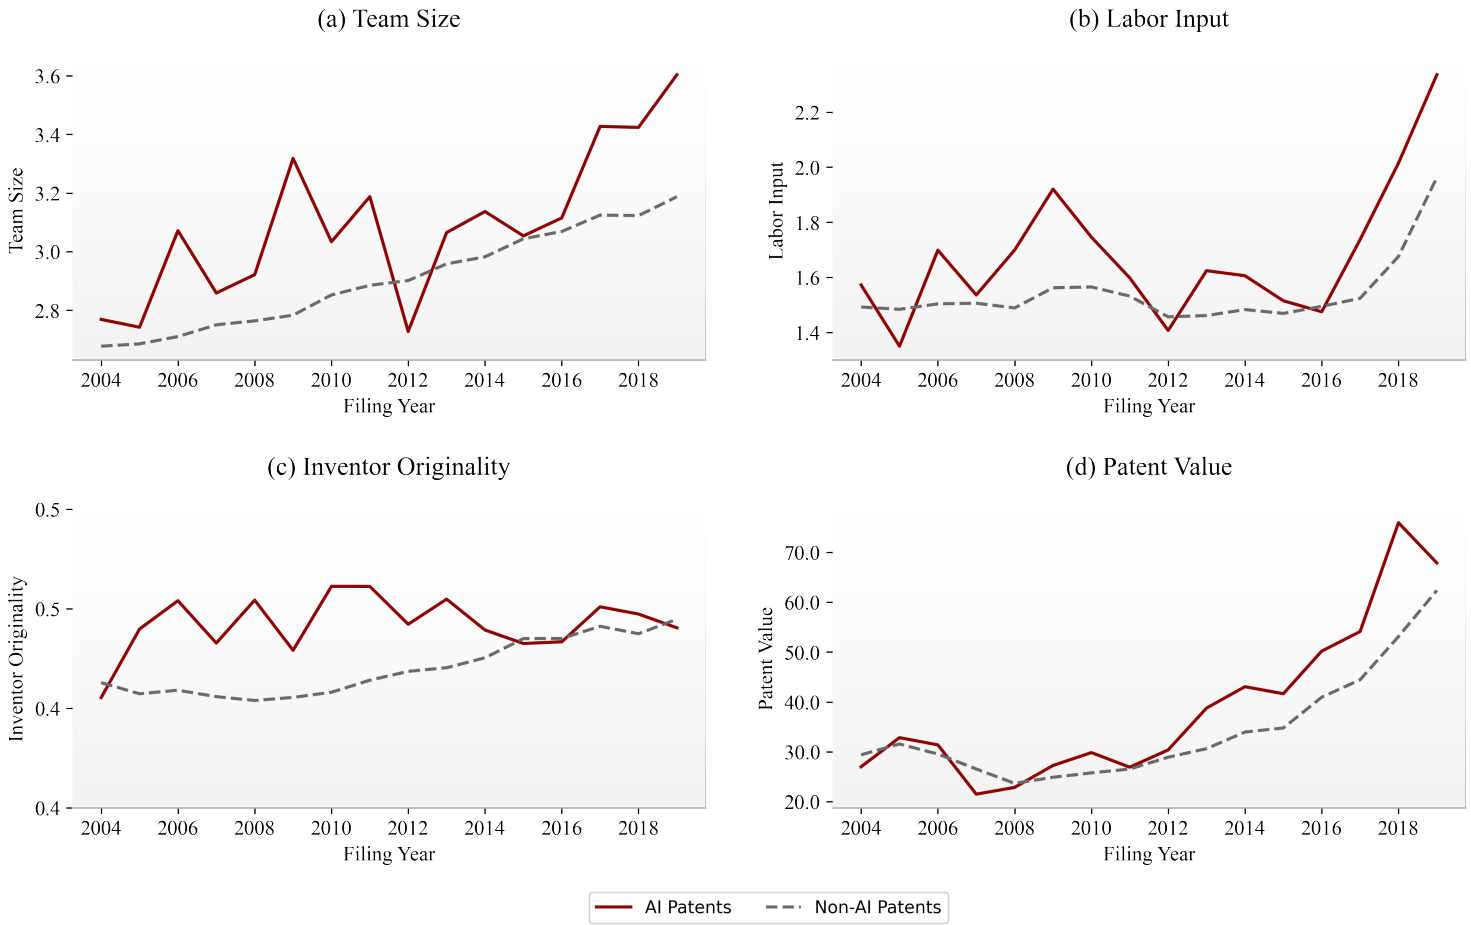


Table I: The Impact of Robot Exposure on Innovation Shift

This table presents the first and second-stage results of the 2SLS estimation examining the relationship between robot exposure and the shift in innovation. Specifically, we estimate:

$$Similarity_{it} = \beta \times Robot\ Exposure_{it} + X_{it}'\gamma + \delta_i + \eta_t + \varepsilon_{it},$$

where *Similarity* is measured using both *Tech Similarity* and *Semantic Similarity*. $Tech\ Similarity_{it}$ is quantified by the cosine similarity of the firm i's patent distributions between years t and $t+5$. $Semantic\ Similarity_{it}$ is calculated using the cosine similarity of the firm i's 768-dimensional continuous vector generated by the PatentBERT model between years t and $t+5$. $Robot\ Exposure_{it}$ measures the degree of robot exposure by firm i in year t and is standardized to have a mean of zero and a standard deviation of one. Columns (1) and (3) present the first-stage results. Columns (2) and (4) report the second-stage estimates of the IV results for the coefficients of robot exposure on technology similarity and semantic similarity for year t, respectively. Firm-level control variables include firm size (measured by total assets), firm age, leverage ratio, return on assets, and proportion of tangible assets. The sample period covers the years 2004–2019, and the regression data are at the firm-year level. All regressions control for firm-fixed effects and year-fixed effects. Firm-level clustered standard errors are reported in parentheses. Coefficients marked with *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Level	Firm-Year			
Period	t			
Model	IV			
Stage	First Stage	Second Stage	First Stage	Second Stage
Variable	Robot Exposure	Tech Similarity	Robot Exposure	Semantic Similarity
	(1)	(2)	(3)	(4)
Robot Exposure IV	0.003***		0.003***	
	(0.001)		(0.001)	
Robot Exposure		-0.300*		-0.436**
		(0.180)		(0.215)
Controls	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	20,815	20,815	18,238	18,238
F Statistics	24.000		18.740	
P-value (Underidentification)	0.000		0.000	
P-value (Endogeneity)	0.075		0.004	

Table II: The Impact of Robot Exposure on AI-Directed Innovation Shift

This table investigates whether firms shift their innovation activities towards AI-related areas in response to increased robot exposure. The regressions in Panel A use firm-year level data, whereas Panel B uses patent-firm-year level data. Panel A estimates the following equation:

$$y_{it} = \beta \times \textit{Robot Exposure}_{it} + X'_{it}\gamma + \delta_i + \eta_t + \varepsilon_{it},$$

where the dependent variable y is measured using both *AI Closeness* and *AI Distance*. For each patent p, $\textit{AI Closeness}_{ipt}$ is defined as the cosine similarity between the citation distribution vector of patent p filed by firm i and that of AI patents in year t. $\textit{AI Distance}_{ipt}$ captures the shortest path length from patent p to any AI patent in the citation network of year t. These measures are aggregated at the firm-year level to construct $\textit{AI Closeness}_{it}$ and $\textit{AI Distance}_{it}$. $\textit{Robot Exposure}_{it}$ denotes the standardized level of robot adoption by firm i in year t, with mean zero and standard deviation one. Columns (1) and (3) report the first-stage results, while Columns (2) and (4) present the second-stage IV estimates of the effect of robot exposure on AI closeness and AI distance, respectively. In Panel B, the dependent variables are the AI closeness of each patent and the logarithm of AI distance (plus one), respectively. Panel C examines the dynamic effects of robot exposure on AI-related innovation in subsequent years. In Panel D, we model AI patenting as a time-varying Poisson process, with the expected arrival rate specified as:

$$\mathrm{E}\left(\mathrm{y}_{\mathrm{i,t+j}} \mid \mathbf{X}_{\mathrm{it}}\right) = \exp\left(\beta \times \textit{Robot Exposure}_{it} + \mathbf{X}'_{it}\gamma + \delta_i + \eta_t\right),$$

where $\mathrm{E}\left(\mathrm{y}_{\mathrm{i,t+j}} \mid \mathbf{X}_{\mathrm{it}}\right)$ represents the expected AI patent counts filed by firm i in year $t+j$. Firm-level control variables include firm size (measured by total assets), firm age, leverage ratio, return on assets, and the proportion of tangible assets. The sample period spans 2004–2019. All regressions incorporate firm-fixed and year-fixed effects. Firm-level clustered standard errors are reported in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Model	IV			
Period	t			
Stage	First Stage	Second Stage	First Stage	Second Stage
Variable	Robot Exposure	AI Closeness	Robot Exposure	AI Distance
	(1)	(2)	(3)	(4)
	Panel A: Firm-Year Level			
Robot Exposure IV	0.003***		0.003***	
	(0.001)		(0.001)	
Robot Exposure		0.004**		-0.736**
		(0.002)		(0.333)
Observations	20,815	20,815	20,815	20,815
F Statistics	24.000		24.000	
P-value (Underidentification)	0.000		0.000	
P-value (Endogeneity)	0.002		0.002	
	Panel B: Patent-Firm-Year Level			
Robot Exposure IV	0.015***		0.015***	
	(0.000)		(0.000)	
Robot Exposure		0.081***		-0.048***
		(0.004)		(0.012)
Observations	607,772	607,772	607,772	607,772
F Statistics	14.460		14.460	
P-value (Underidentification)	0.000		0.000	
P-value (Endogeneity)	0.000		0.000	

(Continued)

Level	Firm-Year					
	Panel C: Long-Term Effects					
Model	IV					
Stage	Second Stage					
Variable	AI Closeness			AI Distance		
Period	t+1	t+2	t+3	t+1	t+2	t+3
	(1)	(2)	(3)	(4)	(5)	(6)
Robot Exposure	0.004*	0.006*	0.004*	-0.799*	-0.602	-0.042
	(0.002)	(0.003)	(0.002)	(0.443)	(0.475)	(0.416)
Observations	18,088	15,908	13,945	18,088	15,908	13,945
F Statistics	15.580	11.550	11.150	15.580	11.550	11.150
P-value (Underidentification)	0.003	0.010	0.010	0.003	0.010	0.010
P-value (Endogeneity)	0.001	0.001	0.001	0.001	0.001	0.001
	Panel D: AI Patent Counts					
Model	Poisson					
Variable	AI Patent Count					
Period	t	t+1	t+2	t+3		
	(1)	(2)	(3)	(4)		
Robot Exposure	0.211***	0.244***	0.170**	0.004		
	(0.066)	(0.077)	(0.085)	(0.099)		
Observations	3,058	2,692	2,396	2,083		
Controls	Yes	Yes	Yes	Yes		
Firm FE	Yes	Yes	Yes	Yes		
Year FE	Yes	Yes	Yes	Yes		

Table III: The Impact of Robot Exposure on 8 Types of AI Closeness/AI Distance

This table presents the second-stage results of the 2SLS estimation, analyzing which specific AI-related fields firms adjust their innovations towards in response to increased robot exposure. We employ an alternative measure of AI patents following H. Fang et al. 2025. Here, we consider eight subfields of AI, including machine learning (ML), evolutionary computation (EVO), natural language processing (NLP), speech processing (SPEECH), computer vision (VISION), knowledge representation (KR), planning and control (PLAN), and AI hardware (HW). We divide each AI patent into one or more of these fields and calculate the closeness/distance for any given patent to the specific field AI patents. Specifically, we estimate:

$$y_{it} = \beta \times Robot\ Exposure_{it} + X'_{it}\gamma + \delta_i + \eta_t + \varepsilon_{it},$$

where the dependent variable y is measured using both *AI Closeness* and *AI Distance*. $Robot\ Exposure_{it}$ measures the degree of robot exposure by firm i in year t and is standardized to have a mean of zero and a standard deviation of one. Firm-level control variables include firm size (measured by total assets), firm age, leverage ratio, return on assets, and the proportion of tangible assets. The sample period spans 2004–2019, and the regression data are at the firm-year level. All regressions incorporate firm-fixed and year-fixed effects. Firm-level clustered standard errors are reported in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Level	Firm-Year							
Model	IV							
Stage	Second Stage							
Period	t							
Types	ML	EVO	NLP	SPEECH	VISION	KR	PLAN	HW
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Alternative AI Patent Measure: H. Fang et al. 2025							
Variable	***Panel A: AI Closeness***							
Robot Exposure	2.617**	2.573**	2.679**	2.609**	2.646**	2.653**	2.573**	2.643**
	(1.162)	(1.145)	(1.193)	(1.151)	(1.193)	(1.166)	(1.174)	(1.187)
Observations	20,815	20,815	20,815	20,815	20,815	20,815	20,815	20,815
F Statistics	24.000	24.000	24.000	24.000	24.000	24.000	24.000	24.000
P-value (Underidentification)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
P-value (Endogeneity)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Variable	***Panel B: AI Distance***							
Robot Exposure	0.036	-0.007	-0.030	-0.015	-0.054	-0.044	0.005	-0.019
	(0.231)	(0.229)	(0.217)	(0.218)	(0.233)	(0.226)	(0.233)	(0.240)
Observations	20,815	20,815	20,815	20,815	20,815	20,815	20,815	20,815
F Statistics	24.000	24.000	24.000	24.000	24.000	24.000	24.000	24.000
P-value (Underidentification)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
P-value (Endogeneity)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table IV: The Impact of Robot Exposure on Data Generation

This table presents the results of the 2SLS estimation, which investigates the relationship between robot exposure and data generation. Specifically, we estimate the following equation:

$$Data\ Generation_{it} = \beta \times Robot\ Exposure_{it} + X'_{it}\gamma + \delta_i + \eta_t + \varepsilon_{it},$$

where $Data\ Generation_{it}$ represents the frequency of data-generation-related keywords in firm i's 10-K report for year t, normalized by the total word count of the report for that year. $Robot\ Exposure_{it}$ measures the degree of robot exposure by firm i in year t. Both *Data Generation* and *Robot Exposure* are standardized to have a mean of zero and a standard deviation of one. Columns (1) and (2) present the first and second stage coefficients of robot exposure on data generation, respectively. Columns (3)-(5) present the results of the heterogeneity analysis for firm size. Large firms are identified using *Dummy* $(Sales > 75\%)$, which equals one if the firm's sales revenue is in the top 25% within its industry and year. Firm-level control variables include firm size (measured by total assets), firm age, leverage ratio, return on assets, and the proportion of tangible assets. The sample period spans the years 2004–2019, and the regression data are at the firm-year level. All regressions control for firm-fixed effects and year-fixed effects. Firm-level clustered standard errors are reported in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Level	Firm-Year				
Period	t				
Model	IV				
Stage	First Stage	Second Stage	First Stage		Second Stage
Variable	Robot Exposure	Data Generation	Robot Exposure	Robot Exposure × Dummy (Sales > 75%)	Data Generation
	(1)	(2)	(3)	(4)	(5)
Robot Exposure IV	0.003***		0.004***	-0.000	
	(0.001)		(0.001)	(0.001)	
Robot Exposure IV × Dummy (Sales > 75%)			-0.000	0.004***	
			(0.001)	(0.001)	
Robot Exposure		1.313***			1.043**
		(0.485)			(0.476)
Robot Exposure × Dummy (Sales > 75%)					0.739*
					(0.439)
Dummy (Sales > 75%)			0.014	0.021	-0.032
			(0.018)	(0.018)	(0.044)
Controls	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	20,815	20,815	20,815	20,815	20,815
F Statistics	25.980			13.060	
P-value (Underidentification)	0.000			0.000	
P-value (Endogeneity)	0.000			0.000	

Table V: The impact of Robot Exposure on the Cumulative Innovation Outputs

This table examines the impact of robot exposure on firms' cumulative innovation outputs over multiple years. Panel A and Panel B present results for the cumulative number of patents and cumulative patent value, respectively. Columns (1)–(6) in each panel report Poisson and second-stage 2SLS estimates over horizons from 1 to 6 years. Patent grants are modeled as a time-varying Poisson process, with the expected arrival rate specified as:

$$\mathrm{E}\left(\mathrm{y_{i,t+j}} \mid \mathbf{X_{it}}\right) = \exp\left(\beta \times \textit{Robot Exposure}_{it} + \mathbf{X}'_{it}\gamma + \delta_i + \eta_t\right),$$

where $\mathrm{E}\left(\mathrm{y_{i,t+j}} \mid \mathbf{X_{it}}\right)$ denotes the expected cumulative number of patents filed by firm i from year t to year $t+j$. The patent count is adjusted following Acemoglu, Akcigit, and Kerr (2016), by normalizing forward five-year citations using the average citation count within each 3-digit CPC class. $\textit{Robot Exposure}_{it}$ captures the extent of robot adoption by firm i in year t, standardized to have mean zero and unit variance. For patent value, we estimate:

$$y_{i,t+j} = \beta \times \textit{Robot Exposure}_{it} + X'_{it}\gamma + \delta_i + \eta_t + \varepsilon_{it},$$

where $y_{i,t+j}$ represents the cumulative patent value filed by firm i from year t to year $t+j$, following the methodology of Kogan et al. (2017). Specifically, we deflate forward five-year citations and aggregate the resulting patent values at the firm level. Firm-level control variables include firm size (measured by total assets), firm age, leverage ratio, return on assets, and the proportion of tangible assets. The sample period spans the years 2004–2019, and the regression data are at the firm-year level. All regressions control for firm-fixed effects and year-fixed effects. Firm-level clustered standard errors are reported in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Level	Firm-Year					
Period	t					
Cumulative	1 year	2 years	3 years	4 years	5 years	6 years
	(1)	(2)	(3)	(4)	(5)	(6)
Model	Poisson					
Variable	***Panel A: Citation-Deflated Patent Count***					
Robot Exposure	-0.086**	-0.061*	-0.042	-0.016	0.018	0.040*
	(0.035)	(0.034)	(0.034)	(0.029)	(0.021)	(0.023)
Observations	16,473	14,603	12,804	11,090	9,536	8,175
Model	IV					
Variable	***Panel B: Citation-Deflated Patent Value***					
Robot Exposure	-0.997**	-1.163**	-1.298*	-1.094*	-0.988	-0.537
	(0.475)	(0.583)	(0.757)	(0.602)	(0.665)	(0.445)
Observations	20,815	18,088	15,630	13,429	11,483	9,801
F Statistics	24.000	15.980	10.160	11.730	8.549	10.260
P-value (Underidentification)	0.000	0.003	0.017	0.009	0.030	0.022
P-value (Endogeneity)	0.004	0.002	0.003	0.003	0.013	0.090
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table VI: The Impact of Robot Exposure on R&D Expenditures

This table presents the results of the 2SLS estimation, which investigates the relationship between robot exposure and R&D expenditures. Specifically, we estimate the following equation:

$$y_{it} = \beta_1 \times Robot\ Exposure_{it} + \beta_2 \times R\&D\ Dummy_{it} + X'_{it}\gamma + \delta_i + \eta_t + \varepsilon_{it},$$

where y_{it} represents firm i's R&D expenditure ratios in year t. We examine two measures of R&D expenditure: R&D normalized by total assets and R&D normalized by total sales. $Robot\ Exposure_{it}$ captures the degree of robot exposure for firm i in year t, standardized to have a mean of zero and a standard deviation of one. $R\&D\ Dummy_{it}$ is a dummy variable to indicate whether the firm i has R&D expenditure in year t. Columns (1) and (3) present the first-stage results. Columns (2) and (4) report the second-stage estimates of the IV results for the coefficients of robot exposure on $R\&D/Total\ Assets$ and $R\&D/Sales$ for year t, respectively. Firm-level control variables include firm size (measured by total assets), firm age, leverage ratio, return on assets, and the proportion of tangible assets. The sample period spans the years 2004–2019, and the regression data are at the firm-year level. All regressions control for firm-fixed effects and year-fixed effects. Firm-level clustered standard errors are reported in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Level	Firm-Year			
Period	t			
Model	IV			
Stage	First Stage	Second Stage	First Stage	Second Stage
Variable	Robot Exposure	R&D/Total Assets	Robot Exposure	R&D/Sales
	(1)	(2)	(3)	(4)
Robot Exposure IV	0.003***		0.003***	
	(0.001)		(0.001)	
Robot Exposure		0.104*		0.385**
		(0.060)		(0.187)
R&D Dummy	0.057**	0.329***	0.001	0.259***
	(0.028)	(0.011)	(0.008)	(0.008)
Controls	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	20,815	20,815	20,815	20,815
F Statistics	23.340		23.980	
P-value (Underidentification)	0.000		0.000	
P-value (Endogeneity)	0.021		0.000	

Table VII: The Impact of Robot Exposure on Patent Generality/Originality

This table presents the two-stage results of the 2SLS estimation, examining how a firm's robot exposure affects its patent generality and originality. The regression data used in Panel A are at the firm-year level, while the data in Panel B are at the patent-firm-year level. Specifically, we estimate the following equation in Panel A:

$$y_{it} = \beta \times Robot\ Exposure_{it} + X'_{it}\gamma + \delta_i + \eta_t + \varepsilon_{it},$$

where y_{it} denotes the aggregated innovation quality of firm i's patents at the firm level in year t. The innovation quality measure includes two components: *Generality* is calculated as one minus the sum of the shares of citations that a given patent receives from any technology class within five years of its issuance. *Originality* is measured as one minus the sum of the shares of patents that a given patent cites across any technology class. Empirically, we define technology classes at the 3-digit CPC level. $Robot\ Exposure_{it}$ measures the degree of robot exposure by firm i in year t and is standardized to have a mean of zero and a standard deviation of one. Column (1) and (3) present the first-stage result. Columns (2) and (4) report the second-stage estimates of the IV results for the coefficients of robot exposure on generality and originality for year t, respectively. Firm-level control variables include firm size (measured by total assets), firm age, leverage ratio, return on assets, and the proportion of tangible assets. The sample period spans 2004–2019. All regressions incorporate firm-fixed and year-fixed effects. Firm-level clustered standard errors are reported in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Period	t			
Model	IV			
Stage	First Stage	Second Stage	First Stage	Second Stage
Variable	Robot Exposure	Generality	Robot Exposure	Originality
	(1)	(2)	(3)	(4)
	Panel A: Firm-Year Level			
Robot Exposure IV	0.003***		0.003***	
	(0.001)		(0.001)	
Robot Exposure		0.135*		0.028**
		(0.070)		(0.013)
Observations	20,815	20,815	20,815	20,815
F Statistics	24.000		24.000	
P-value (Underidentification)	0.000		0.000	
P-value (Endogeneity)	0.017		0.003	
	Panel B: Patent-Firm-Year Level			
Robot Exposure IV	0.016***		0.017***	
	(0.000)		(0.000)	
Robot Exposure		0.042***		0.041***
		(0.005)		(0.004)
Observations	365,009	365,009	607,772	607,772
F Statistics	21.550		19.900	
P-value (Underidentification)	0.000		0.000	
P-value (Endogeneity)	0.000		0.000	
Controls	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes

Table VIII: The Effect of Prior AI Research Experience on AI-Directed Innovation Shift

This table presents the two-stage results of the 2SLS estimation, examining the relationship between firms' prior AI research experience and their shift towards AI-directed innovation at the firm-year level. Specifically, we estimate the following equation:

$$y_{it} = \beta_1 \times Dummy(AI\ Experience)_{i,03} \times Robot\ Exposure_{it} + \beta_2 \times Robot\ Exposure_{it} + X'_{it}\gamma + \delta_i + \eta_t + \varepsilon_{it},$$

where y_{it} is the *AI Closeness* (*Distance*) of firm i at year t. For each patent p, $AI\ Closeness_{ipt}$ is calculated as the cosine similarity between the citation distribution vectors of firm i's patent p and AI patents in year t. $AI\ Distance_{ipt}$ represents the minimum number of steps required for firm i's patent p to reach AI patents within the citation network in year t. These two metrics are aggregated at the firm level. $Robot\ Exposure_{it}$ measures the degree of robot exposure by firm i in year t. $Dummy(AI\ Experience)_{i,03}$ is a dummy variable equal to 1 if firm i's AI experience prior to 2004 is above the industry median for that year, and 0 otherwise. Firm-level control variables include firm size (measured by total assets), firm age, leverage ratio, return on assets, and the proportion of tangible assets. The sample period spans the years 2004–2019, and the regression data are at the firm-year level. All regressions control for firm-fixed effects and year-fixed effects. Firm-level clustered standard errors are reported in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

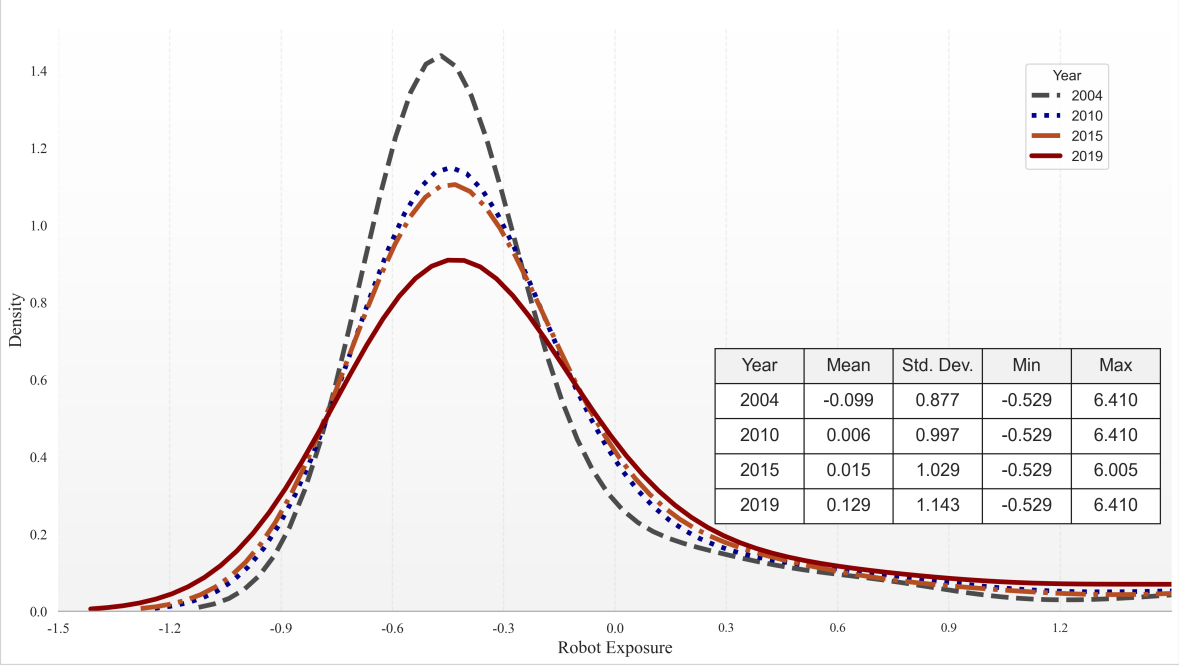
Level	Firm-Year					
Period	t					
Model	IV					
Stage	First Stage		Second Stage	First Stage		Second Stage
Variable	Robot Exposure	Robot Exposure × Dummy (AI Exp)	AI Closeness	Robot Exposure	Robot Exposure × Dummy (AI Exp)	AI Distance
	(1)	(2)	(3)	(4)	(5)	(6)
Robot Exposure IV	0.003***	-0.000		0.003***	-0.000	
	(0.001)	(0.000)		(0.001)	(0.000)	
Robot Exposure IV × Dummy (AI Exp)	0.004	0.010***		0.004	0.010***	
	(0.003)	(0.003)		(0.003)	(0.003)	
Robot Exposure			0.001			-0.768*
			(0.001)			(0.396)
Robot Exposure × Dummy (AI Exp)			0.011**			0.123
			(0.005)			(0.388)
Observations	20,815	20,815	20,815	20,815	20,815	20,815
F Statistics		10.280			10.280	
P-value (Underidentification)		0.001			0.001	
P-value (Endogeneity)		0.000			0.000	
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Internet Appendix

Appendix A: Figures

Figure A1: Distribution of Robot Exposure

This figure shows the distribution of robot exposure in the years 2004, 2010, 2015, and 2019. Robot exposure is defined as the frequency of robot-related keywords in 10-K filings, normalized by the total word count in each report and standardized to have a mean of 0 and a variance of 1. The figure also presents summary statistics for robot exposure in these four years, including the mean, standard deviation, minimum, and maximum values.



Year	Mean	Std. Dev.	Min	Max
2004	-0.099	0.877	-0.529	6.410
2010	0.006	0.997	-0.529	6.410
2015	0.015	1.029	-0.529	6.005
2019	0.129	1.143	-0.529	6.410

Figure A2: Distribution of Incremental Changes in Robot Exposure for the Top 10% of Firms

This figure illustrates the incremental changes in robot exposure across four key years: 2006, 2010, 2015, and 2019. For each of these years, we identify the top 10% of firms based on the magnitude of the change in robot exposure from the previous year. We then compute the robot exposure change in the subsequent year for these firms and present the distribution of these incremental changes.

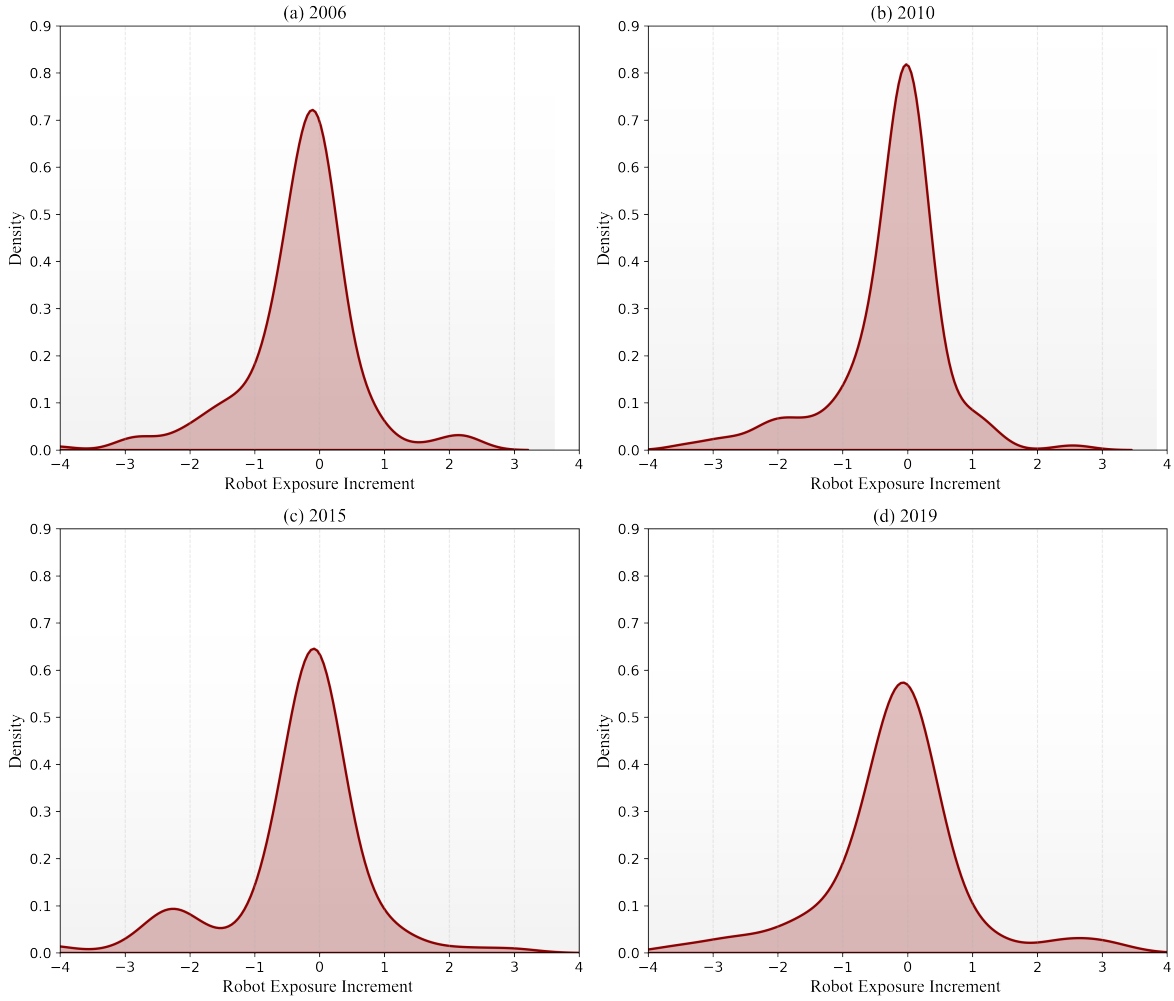


Figure A3: Difference in AI Closeness between AI Patents and Non-AI Patents

This figure compares the difference in AI closeness between AI patents and non-AI patents. Specifically, we randomly designate 10% of AI patents as non-AI patents while treating the remaining 90% as "representative" AI patents. We then compute the AI closeness of all other non-AI patents, including both the sampled pseudo non-AI patents and genuine non-AI patents, relative to the "representative" AI patents. The resulting AI closeness metrics are presented in histograms to compare the distribution of AI closeness between the pseudo non-AI patents and genuine non-AI patents.

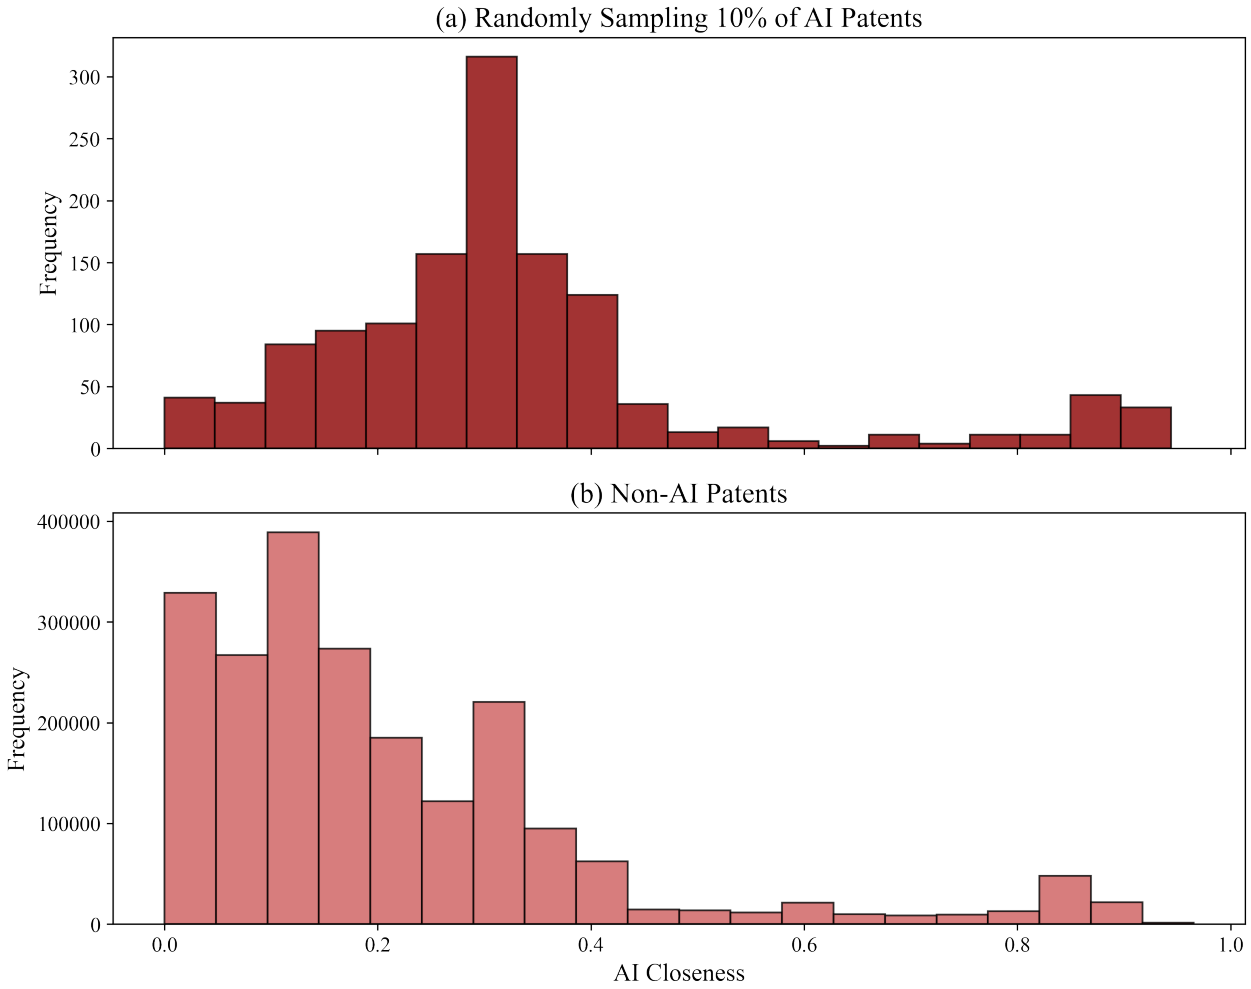
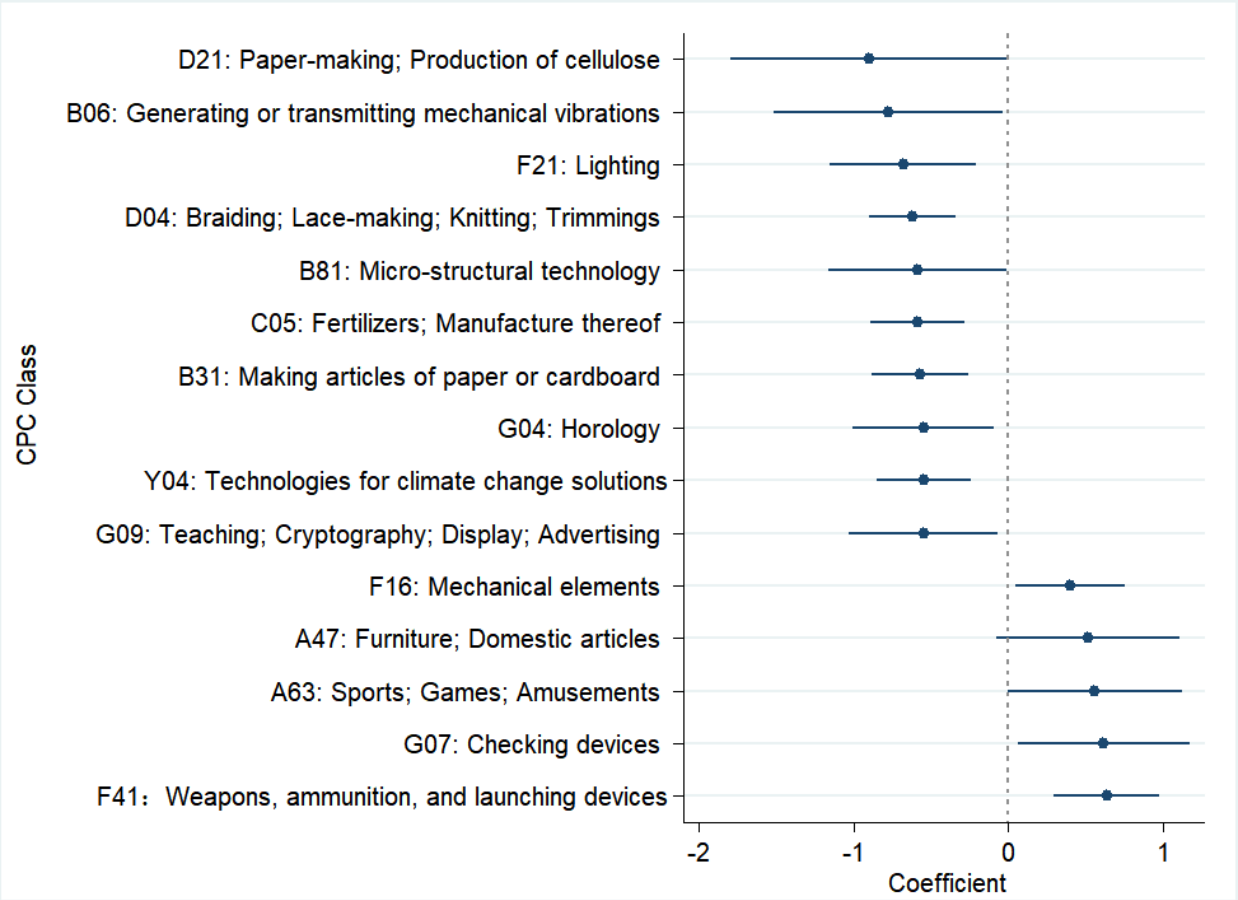


Figure A4: The Impact of Robot Exposure on Different Types of Patent Classes

This figure examines how robot exposure affects different patent classes, categorized by two-digit CPC codes. To capture heterogeneity, we estimate a Poisson model with interaction terms between robot exposure and CPC category indicators:

$$\mathbb{E}\left(y_{i,t+j} \mid \mathbf{X}_{it}\right) = \exp\left(\sum_{n} \rho_n \left(CPC_n \times Robot\ Exposure_{it}\right) + \beta \times Robot\ Exposure_{it} + \mathbf{X}'_{it}\,\gamma + \delta_i + \eta_t\right),$$

where $y_{i,t+j}$ is the expected number of patents filed by firm i in year $t+j$. Each CPC_n is an indicator for patents classified under CPC class n. *Robot Exposure* measures the degree of robot exposure by firm i in year t. Firm-level control variables include firm size (measured by total assets), firm age, leverage ratio, return on assets, and the proportion of tangible assets. The model also includes firm-fixed effects and year-fixed effects. The sample period spans 2004-2019. The figure highlights the top ten CPC classes with statistically significant interaction coefficients, indicating whether robot exposure amplifies or dampens innovation in those specific technology classes.



Appendix B: Tables

Table A1: Summary of the Match between Compustat and USPTO

This table summarizes the matching process between Compustat and USPTO data from 2004 to 2019. For each year, the table lists the total number of Compustat firms, the number (and percentage) of those firms that have at least one patent, and the total number of matched patents granted to these firms.

Year	# of Firms	# of Firms with at least one patent	% of Firms with at least one patent	# of Matched patents in our sample
2004	1,717	853	49.68%	39,095
2005	1,690	829	49.05%	40,770
2006	1,629	823	50.52%	39,914
2007	1,571	808	51.43%	41,283
2008	1,416	735	51.91%	41,819
2009	1,353	720	53.22%	34,222
2010	1,310	733	55.95%	38,425
2011	1,265	748	59.13%	43,131
2012	1,243	748	60.18%	52,235
2013	1,248	774	62.02%	52,859
2014	1,220	713	58.44%	48,754
2015	1,186	695	58.60%	51,868
2016	1,130	657	58.14%	50,442
2017	1,079	628	58.20%	45,509
2018	1,029	572	55.59%	29,869
2019	876	485	55.37%	11,045
Total	20,962	11,521	54.96%	661,240

Table A2: Keyword List for Defining Robot Exposure

This table presents the keywords used to define firm-level robot exposure, based on a comprehensive review of literature on robot adoption. A substantial portion of these keywords derives from IFR Reports, which provide detailed definitions of industrial robots and their applications. When counting keyword frequencies, case sensitivity and plural/singular distinctions are disregarded.

References	Keywords		
IFR Reports	Adopt robot	More robot	Robot invest
Acemoglu and Restrepo (2019)	Assemble robot	Multiple robot	Robot labor
Acemoglu and Restrepo (2020)	Automatically controlled manipulator	Multipurpose manipulator	Robot manufacturing
Acemoglu and Restrepo (2022)	Automation application	Program robot	Robot module
Berg, Buffie, and Zanna (2018)	Automated machine	Reprogrammable manipulator	Robot motor
Caselli and Manning (2019)	Automated robot	Rise of robot	Robot revolution
Dixon, Hong, and L. Wu (2021)	Automated system	Robot adopt	Robot stock
Graetz and Michaels (2018)	Automated technology	Robot advance	Robot supply
Guerreiro, Rebelo, and Teles (2022)	Collaborative robot	Robot application	Robot supporting
Zeira (1998)	Collaborative system	Robot arithmetic	Robot system
	Computer-assisted machine	Robot arm	Robot technology
	Co-robot	Robot automation	Robot utilization
	Employ robot	Robot capacity	Robot vehicle
	Exposure to robot	Robot capital	Robot welding
	Flat panel handler	Robot control	Robot-based
	Full-automated	Robot decelerator	Robot-driven
	Full-process robot	Robot dedicated	Robotic arm
	Handling operation	Robot density	Robotic hardware
	IFR	Robot employ	Robotic platform
	Industrial machine	Robot exposure	Robotic software
	Industrial robot	Robot gripper	Robot-related
	Install robot	Robot head	Robot-specialized
	Labor-replacing machine	Robot input	Robot-to-labor
	Machine tending	Robot install	Smart machine
	Maintain robot	Robot intensity	Stock of robot
	Manufacture robot	Robot introduction	Use of robot

Table A3: Correspondence of Industry Classifications

This table presents the mapping of industries based on 2-digit SIC codes, which are aggregated into 15 broad categories (SIC15), and shows the correspondence between SIC15 categories, IFR industry codes, and EU KLEMS industry classifications.

SIC2	SIC15	SIC15 Industry Description	IFR Industry	EU KLEMS Industry
01–09	1	Agriculture, hunting, forestry, and fishing	A–B	A
10–14	2	Mining and quarrying	C	B
20–21	3	Food products and beverages; tobacco products	D10–D12	C10–C12
22–23, 31	4	Textiles, leather, wearing apparel	D13–D15	C13–C15
24–25, 26–27	5	Wood and wood products, paper and paper products, publishing and printing	D16–D18	C16–C18
28–30, 32	6	Plastic and chemical products, glass, ceramics, stone, mineral products	D19–D23	C19–C23
33–34	7	Basic metals and metal products	D24–D25	C24–C25
35	8	Metal machinery	D28	C28
36, 38	9	Electrical/electronics	D26–D27	C26–C27
37	10	Automotive and other transport equipment	D29–D30	C29–C30
39	11	All other manufacturing branches	91	C31–C33
49	12	Electricity, gas and water supply	E	D–E
15–17	13	Construction	F	F
82, 87	14	Education, research and development	P	M–Q
40–48, 50–69, 70–81, 83–86, 88–99	15	All other non-manufacturing branches	90	G–L, R–U

Table A4: Robustness Checks: Specifications with and without Controls

This table presents robustness results across specifications with and without controls. The dependent variables include: *Tech Similarity*, *Semantic Similarity*, *AI Closeness*, *AI Distance*, *R&D/Total Assets*, *R&D/Sales*, *Patent Value*, and *Patent Count*. *Robot Exposure* quantifies the degree of automation faced by each firm in a given year, standardized to have a mean of zero and a standard deviation of one. Columns (1)–(7) present second-stage results from 2SLS regressions, while Column (8) reports estimates from a Poisson regression. The sample period spans 2004–2019. The regression data are at the firm-year level. All regressions incorporate firm-fixed and year-fixed effects. Firm-level clustered standard errors are reported in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Period	t							
Model	IV							Poission
Variable	Tech Similarity	Semantic Similarity	AI Closeness	AI Distance	R&D/Total Assets	R&D/ Sales	Patent Value	Patent Count
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Panel A: Specification without Controls							
Robot Exposure	-0.198	-0.388**	0.004**	-0.874***	0.118*	0.429**	-1.081**	-0.093**
	(0.157)	(0.183)	(0.002)	(0.332)	(0.064)	(0.182)	(0.448)	(0.040)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
F Statistics	28.180	28.180	28.180	23.180	28.180	28.180	28.180	
P-value (Underidentification)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	
P-value (Endogeneity)	0.004	0.001	0.000	0.002	0.002	0.001	0.004	
Observations	20,815	18,238	20,815	20,815	20,815	20,815	20,815	16,473
	Panel B: Specification with Controls							
Robot Exposure	-0.300*	-0.436**	0.004**	-0.736**	0.104*	0.385**	-0.997**	-0.086**
	(0.180)	(0.215)	(0.002)	(0.333)	(0.060)	(0.187)	(0.475)	(0.035)
Age	-0.115***	-0.062*	-0.000	0.168***	0.061***	0.037	0.087	0.031
	(0.028)	(0.033)	(0.000)	(0.054)	(0.011)	(0.033)	(0.074)	(0.192)
Size	0.016**	0.026***	0.001***	-0.035**	-0.000	0.005	0.262***	0.397***
	(0.008)	(0.008)	(0.000)	(0.015)	(0.003)	(0.011)	(0.022)	(0.082)
ROA	-0.023*	0.014	0.000	0.017	0.020***	0.014	0.119***	-0.129
	(0.013)	(0.014)	(0.000)	(0.021)	(0.005)	(0.016)	(0.025)	(0.168)
Leverage	-0.008	-0.016	-0.000	0.048**	0.024***	-0.020	-0.057*	-0.260
	(0.014)	(0.015)	(0.000)	(0.024)	(0.006)	(0.013)	(0.032)	(0.199)
Tangibility	-0.079**	-0.016	-0.001***	0.053	0.086***	0.094**	-0.095	-1.397**
	(0.032)	(0.034)	(0.000)	(0.062)	(0.011)	(0.040)	(0.078)	(0.585)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
F Statistics	24.000	18.740	24.000	24.000	24.000	24.000	24.000	
P-value (Underidentification)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	
P-value (Endogeneity)	0.075	0.004	0.002	0.002	0.021	0.000	0.004	
Observations	20,815	18,238	20,815	20,815	20,815	20,815	20,815	16,473

Table A5: Top 20 Words: Relative Frequency Differences between High- and Low-Robot-Exposure Patents

This table displays the top 20 n-grams with the highest relative frequency differences between patents with robot exposure above the median and those below the median, based on samples from 2010, 2015, and 2019. To compute these values, we first normalize n-gram frequencies for each group and year relative to the baseline year of 2004. We then calculate the ratio of normalized frequencies between the high- and low-exposure groups and apply a natural log transformation to obtain group-level log-ratios of n-gram prominence. Figure II presents word clouds corresponding to the results shown in the table.

Year	2010		2015		2019	
No.	N-grams	Freq.	N-grams	Freq.	N-grams	Freq.
1	Relates genetic	4.757	Turbine engine	5.914	Golf club	4.228
2	Still provided	4.670	Article footwear	5.651	Club head	4.228
3	Provided invention	4.670	Rotor blade	5.433	First electromagnetic	4.198
4	Plant plant	4.418	Wind turbine	5.291	Golf ball	4.153
5	Invention thus	4.412	Treatment fluid	5.279	Transmission medium	4.121
6	Plants plant	4.400	Sole structure	5.055	Syntax elements	4.054
7	Automated banking	4.400	Golf club	5.025	Work function	4.054
8	Provided seed	4.400	Gas turbine	4.952	Electromagnetic wave	3.865
9	Banking machine	4.351	Club head	4.911	Function metal	3.865
10	Plants produced	4.339	Drill bit	4.903	Fin structure	3.755
11	Transmitting devices	4.266	Glass sheet	4.841	Disclosure may	3.737
12	Derived use	4.203	Drill string	4.813	Communications described	3.731
13	Parts produced	4.203	Glass article	4.724	Electrical return	3.731
14	Graphical program	4.144	Suction side	4.682	Electromagnetic waves	3.713
15	Transgenic soybean	4.144	Ceramic matrix	4.627	Network controller	3.657
16	Dual clutch	4.144	Drilling fluid	4.615	Requiring electrical	3.605
17	Replacement gate	4.096	Matrix composite	4.592	Core comprising	3.549
18	Gate conductor	4.096	Downhole tool	4.544	Surface transmission	3.549
19	Interface object	4.080	Turbine component	4.531	Numerous aspects	3.549
20	Storage environment	4.080	Nonhuman animals	4.383	Recording tape	3.520

Table A6: Keyword List for Defining AI Patents

This table lists the AI-related keywords used to identify AI patents. In the analysis, we do not distinguish between case sensitivity and singular and plural forms of these keywords.

Keywords	
Bayesian belief network	Neural network
Collaborative system	Neuromorphic computing
Computer vision	Optimal research
Crowdsourcing and human computation	Pattern analysis
Decision-making	Pattern recognition
Deep learning	Physical symbol system
Generative model	Reinforcement learning
Image alignment	Sensor data fusion
Image grammar	Sensor network
Image matching	Supervised learning
Knowledge representation and reasoning	Symbol processing
Layered control system	Symbolic error analysis
Logic theorist	Symbolic reasoning
Machine intelligence	Symbolic system
Machine learning	System and control theory
Natural language processing	Unsupervised learning

Table A7: The Impact of Robot Exposure on 8 Types of AI Closeness/AI Distance (USPTO-Based Measure)

This table presents the second-stage results of the 2SLS estimation, analyzing which specific AI-related fields firms adjust their innovations towards in response to increased robot exposure. We employ an alternative measure of AI patents based on USPTO records. Here, we consider eight subfields of AI, including machine learning (ML), evolutionary computation (EVO), natural language processing (NLP), speech processing (SPEECH), computer vision (VISION), knowledge representation (KR), planning and control (PLAN), and AI hardware (HW). We divide each AI patent into one or more of these fields and calculate the closeness/distance for any given patent to the specific field AI patents. Specifically, we estimate:

$$y_{it} = \beta \times Robot\ Exposure_{it} + X'_{it}\gamma + \delta_i + \eta_t + \varepsilon_{it},$$

where the dependent variable y is measured using both *AI Closeness* and *AI Distance*. $Robot\ Exposure_{it}$ measures the degree of robot exposure by firm i in year t and is standardized to have a mean of zero and a standard deviation of one. Firm-level control variables include firm size (measured by total assets), firm age, leverage ratio, return on assets, and the proportion of tangible assets. The sample period spans 2004–2019, and the regression data are at the firm-year level. All regressions incorporate firm-fixed and year-fixed effects. Firm-level clustered standard errors are reported in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Level	Firm-Year							
Model	IV							
Stage	Second Stage							
Period	t							
Types	ML	EVO	NLP	SPEECH	VISION	KR	PLAN	HW
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Alternative AI Patent Measure: USPTO							
Variable	***Panel A: AI Closeness***							
Robot Exposure	2.908**	2.747**	2.602**	2.665**	2.936**	3.129**	3.364**	2.667**
	(1.294)	(1.223)	(1.158)	(1.180)	(1.315)	(1.429)	(1.519)	(1.191)
Observations	20,815	20,815	20,815	20,815	20,815	20,815	20,815	20,815
F Statistics	24.000	24.000	24.000	24.000	24.000	24.000	24.000	24.000
P-value (Underidentification)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
P-value (Endogeneity)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Variable	***Panel B: AI Distance***							
Robot Exposure	-0.090	-0.128	-0.231	-0.155	-0.104	0.009	-0.110	-0.106
	(2.909)	(2.894)	(2.892)	(2.894)	(2.915)	(2.959)	(2.945)	(2.932)
Observations	20,815	20,815	20,815	20,815	20,815	20,815	20,815	20,815
F Statistics	24.000	24.000	24.000	24.000	24.000	24.000	24.000	24.000
P-value (Underidentification)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
P-value (Endogeneity)	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table A8: Keyword List for Defining Data Generation

This table presents the set of data-generation-related keywords employed to construct firm-level data generation measures. We proxy data generation by counting the occurrence of these keywords in 10-K filings. The first column references a broad range of literature, while the subsequent columns list the keywords. In counting keyword frequencies, we disregard case sensitivity and singular/plural distinctions.

References	Keywords	
Abis and Veldkamp (2024)	Big data	Data sharing
Bajari et al. (2019)	Business data	Data storage
Brynjolfsson and McElheran (2016)	Data accumulation	Data volume
L. Wu, Hitt, and Lou (2020)	Data asset	Database management
C. Zhu (2019)	Data available	Digital technology
	Data capability	Electronic data
	Data collection	Generate data
	Data creation	Information based
	Data driven	Information provision
	Data extraction	Intangible asset
	Data gathering	Intensive information
	Data generation	Massive information
	Data infrastructure	More data
	Data mart	New information
	Data mining	Rich data
	Data process	Rich information

Table A9: Prior AI Research Experience and AI-Directed Innovation Shift (Patent–Firm–Year Analysis)

This table presents the two-stage results of the 2SLS estimation, examining the relationship between firms' prior AI research experience and their shift towards AI-directed innovation at the patent-firm-year level. For each patent p, $AI\ Closeness_{ipt}$ is calculated as the cosine similarity between the citation distribution vectors of firm i's patent p and AI patents in year t. $AI\ Distance_{ipt}$ represents the minimum number of steps required for firm i's patent p to reach AI patents within the citation network in year t. $Robot\ Exposure_{it}$ measures the degree of robot exposure by firm i in year t. $Dummy(AI\ Experience)_{i,03}$ is a dummy variable equal to 1 if firm i's AI experience prior to 2004 is above the industry median for that year, and 0 otherwise. Firm-level control variables include firm size (measured by total assets), firm age, leverage ratio, return on assets, and the proportion of tangible assets. The sample period spans the years 2004–2019. All regressions control for firm-fixed effects and year-fixed effects. Firm-level clustered standard errors are reported in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Level	Patent-Firm-Year					
Period	t					
Model	IV					
Stage	First Stage		Second Stage	First Stage		Second Stage
Variable	Robot Exposure	Robot Exposure × Dummy (AI Exp)	AI Closeness	Robot Exposure	Robot Exposure × Dummy (AI Exp)	AI Distance
	(1)	(2)	(3)	(4)	(5)	(6)
Robot Exposure IV	0.014***	-0.004***		0.014***	-0.004***	
	(0.000)	(0.000)		(0.000)	(0.000)	
Robot Exposure IV × Dummy (AI Exp)	0.001***	0.018***		0.001***	0.018***	
	(0.000)	(0.000)		(0.000)	(0.000)	
Robot Exposure			0.050***			-0.040***
			(0.004)			(0.015)
Robot Exposure × Dummy (AI Exp)			0.053***			-0.049**
			(0.005)			(0.019)
Observations	607,772	607,772	607,772	607,772	607,772	607,772
F Statistics		14.992			16.094	
P-value (Underidentification)		0.000			0.000	
P-value (Endogeneity)		0.000			0.000	
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table A10: The Impact of Robot Exposure on the Cumulative Number of Patents

This table examines the impact of robot exposure on firms' cumulative number of patents over multiple years. Panel A and Panel B present results based on a balanced panel using Poisson regressions and second-stage 2SLS regressions over horizons from 1 to 6 years. Firm-level control variables include firm size (measured by total assets), firm age, leverage ratio, return on assets, and the proportion of tangible assets. The sample period spans 2004–2019, and the data are at the firm-year level. All regressions include firm fixed effects and year fixed effects. Standard errors are clustered at the firm level and reported in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Level	Firm-Year					
Period	t					
Variable	Citation-Deflated Patent Count					
Cumulative	1 year	2 years	3 years	4 years	5 years	6 years
	(1)	(2)	(3)	(4)	(5)	(6)
Model	Poisson					
Sample	Balanced Panel					
Robot Exposure	0.015	0.039	0.044**	0.042*	0.041*	0.040*
	(0.027)	(0.026)	(0.022)	(0.022)	(0.023)	(0.023)
Observations	8,175	8,175	8,175	8,175	8,175	8,175
Model	IV					
Robot Exposure	-0.420*	-0.560*	-0.699	-0.725*	-0.907	-0.685
	(0.226)	(0.330)	(0.457)	(0.414)	(0.551)	(0.419)
Observations	20,815	18,088	15,630	13,429	11,483	9,801
F Statistics	24.000	15.980	10.160	11.730	8.549	10.260
P-value (Underidentification)	0.000	0.003	0.017	0.009	0.030	0.022
P-value (Endogeneity)	0.001	0.001	0.000	0.003	0.001	0.006
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table A11: Comparison of Relative Costs Between AI Patents and Non-AI Patents

This table reports results from an OLS regression analyzing the relative costs associated with AI patents compared to non-AI patents, corresponding to Figure III. Specifically, we estimate the following model:

$$y_{ipt} = \beta \times Dummy(AI\ Patent)_{ipt} + \theta_{it} + \varepsilon_{ipt},$$

where y_{ipt} denotes the cost of patent p filed by firm i in year t. The dependent variable is examined through four distinct measures: *Team Size* is defined as the number of inventors contributing to a patent; *Labor Input* adjusts *Team Size* by addressing any potential double-counting of inventors; *Inventor Originality* is measured as one minus the sum of the shares of patents filed by the inventor in any given technological class; *Patent Value* follows the approach proposed by Kogan et al. (2017). The key independent variable is $Dummy(AI\ patent)$, which takes a value of 1 if the patent is categorized as an AI patent and 0 otherwise. Firm-level control variables include firm size (measured by total assets), firm age, leverage ratio, return on assets, and the proportion of tangible assets. The sample period spans 2004–2019, and the regression data are at the patent-firm-year level. Firm-level clustered standard errors are reported in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Level	Patent-Firm-Year							
Model	OLS							
Period	t							
Variable	Team Size		Labor Input		Inventor Originality		Patent Value	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dummy (AI Patent)	0.118***	0.118***	0.077***	0.081***	0.041***	0.037***	2.975***	1.675***
	(0.025)	(0.025)	(0.018)	(0.018)	(0.003)	(0.003)	(0.538)	(0.481)
Controls	Yes	No	Yes	No	Yes	No	Yes	No
Firm FE	Yes	No	Yes	No	Yes	No	Yes	No
Year FE	Yes	No	Yes	No	Yes	No	Yes	No
Firm×Year FE	No	Yes	No	Yes	No	Yes	No	Yes
Observations	607,772	607,772	607,772	607,772	607,772	607,772	607,772	607,772

Table A12: Robustness Checks: Controlling for Common Shock Effects

This table presents 2SLS robustness results, with Panel A controlling for the industry-level Producer Price Index and Panel B controlling for industry-year fixed effects. For ease of presentation, the Producer Price Index is scaled down by a factor of 1,000. The dependent variables include: *Tech Similarity*, *Semantic Similarity*, *AI Closeness*, *AI Distance*, *R&D/Total Assets*, *R&D/Sales*, and *Patent Value*. *Robot Exposure* quantifies the degree of automation faced by each firm in a given year, standardized to have a mean of zero and a standard deviation of one. Firm-level control variables include firm size (measured by total assets), firm age, leverage ratio, return on assets, and the proportion of tangible assets. The sample period spans 2004–2019. The regression data are at the firm-year level. All regressions incorporate firm-fixed and year-fixed effects. Firm-level clustered standard errors are reported in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Period	t						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Panel A: Controlling for Industry-level PPI						
Model & Stage	***(a) IV: Second Stage***						
Variable	Tech Similarity	Semantic Similarity	AI Closeness	AI Distance	R&D/Total Assets	R&D/ Sales	Patent Value
Robot Exposure	-0.169	-0.428**	0.005**	-0.615**	0.184**	0.447**	-1.037**
	(0.169)	(0.218)	(0.002)	(0.312)	(0.079)	(0.200)	(0.484)
PPI	-0.886***	-0.395**	-0.005***	-1.552***	-0.163***	0.521**	-0.242
	(0.167)	(0.194)	(0.002)	(0.305)	(0.061)	(0.213)	(0.460)
Model & Stage	***(b) IV: First Stage***						
Variable	Robot Exposure						
Robot Exposure IV	0.003***	0.003***	0.003***	0.003***	0.003***	0.003***	0.003***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
PPI	-0.376**	-0.454**	-0.376**	-0.376**	-0.376**	-0.376**	-0.376**
	(0.164)	(0.178)	(0.164)	(0.164)	(0.164)	(0.164)	(0.164)
F Statistics	24.130	17.960	24.130	24.130	24.130	24.130	24.130
P-value (Underidentification)	0.000	0.000	0.000	0.000	0.000	0.000	0.000
P-value (Endogeneity)	0.349	0.007	0.000	0.009	0.002	0.001	0.003
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	20,251	17,856	20,251	20,251	20,251	20,251	20,251
	Panel B: Controlling for Industry-Year Fixed Effects						
Model & Stage	***(a) IV: Second Stage***						
Variable	Tech Similarity	Semantic Similarity	AI Closeness	AI Distance	R&D/Total Assets	R&D/ Sales	Patent Value
Robot Exposure	-0.009	-0.421	0.007*	-0.836	0.130	0.689*	-0.576
	(0.271)	(0.312)	(0.004)	(0.571)	(0.114)	(0.405)	(0.667)
Model & Stage	***(b) IV: First Stage***						
Variable	Robot Exposure						
Robot Exposure IV	0.002**	0.002**	0.002**	0.002**	0.002**	0.002**	0.002**
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
F Statistics	9.512	8.955	9.512	9.512	9.512	9.512	9.512
P-value (Underidentification)	0.024	0.022	0.024	0.024	0.024	0.024	0.024
P-value (Endogeneity)	0.964	0.061	0.000	0.028	0.179	0.001	0.294
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
SIC2×Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	20,803	18,220	20,803	20,803	20,803	20,803	20,803

Table A13: Robustness Checks: Pre-trend Tests

This table presents robustness check results examining whether there are any pre-existing trends in firms' prior innovation activities. In Panel A, we assess the effect of robot exposure during the 2004–2019 period on firms' innovation shift and innovation productivity between 1988 and 2003. In Panel B, we directly control for changes in firms' innovation activities between 1988 and 2003. The dependent variables include: *Tech Similarity*, *Semantic Similarity*, *AI Closeness*, *AI Distance*, *R&D/Total Assets*, *R&D/Sales*, and *Patent Value*. *Robot Exposure* quantifies the degree of automation faced by each firm in a given year, standardized to have a mean of zero and a standard deviation of one. Firm-level control variables include firm size (measured by total assets), firm age, leverage ratio, return on assets, and the proportion of tangible assets. The sample period spans from 2004 to 2019. The regression data for all panels are structured at the firm-year level. All regressions incorporate firm-fixed and year-fixed effects. Firm-level clustered standard errors are reported in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Model	IV						
Period	t						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Panel A: Pre-trend Tests						
Stage	Second Stage						
Variable Period	1988–2003						
Variable	Tech Similarity	Semantic Similarity	AI Closeness	AI Distance	R&D/Total Assets	R&D/ Sales	Patent Value
Robot Exposure (2004–2019)	-0.484	-0.380	-0.003	0.764	0.554	-0.088	0.967
	(0.749)	(1.795)	(0.005)	(0.915)	(0.604)	(0.251)	(1.690)
Observations	10,878	3,298	10,878	10,878	10,878	10,878	10,878
	Panel B: Controlling for the Change in the Dependent Variable						
Stage	***(a) Second Stage***						
Variable Period	2004–2019						
Variable	Tech Similarity	Semantic Similarity	AI Closeness	AI Distance	R&D/Total Assets	R&D/ Sales	Patent Value
Robot Exposure (2004–2019)	-0.298*	-0.435**	0.004**	-0.736**	0.181**	0.473**	-0.996**
	(0.179)	(0.214)	(0.002)	(0.333)	(0.079)	(0.204)	(0.475)
Change in Dep. Var. (1988–2003)	-0.020**	-0.015	0.006	-0.008	-0.003	0.002	-0.001
	(0.010)	(0.020)	(0.061)	(0.013)	(0.013)	(0.009)	(0.015)
Observations	20,815	18,238	20,815	20,815	20,815	20,815	20,815
Stage	***(b) First Stage***						
Variable Period	2004–2019						
Variable	Robot Exposure						
Robot Exposure IV (2004–2019)	0.003***	0.003***	0.003***	0.003***	0.003***	0.003***	0.003***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Change in Dep. Var. (1988–2003)	-0.008	0.001	2.677	-0.004	-0.046	-0.007	0.001
	(0.011)	(0.025)	(3.005)	(0.008)	(0.035)	(0.008)	(0.006)
Observations	20,815	18,238	20,815	20,815	20,815	20,815	20,815
F Statistics	24.030	18.740	23.960	24.000	23.850	23.940	23.990
P-value (Underidentification)	0.000	0.000	0.000	0.000	0.000	0.000	0.000
P-value (Endogeneity)	0.077	0.005	0.002	0.002	0.003	0.001	0.005
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table A14: Robustness Checks: Alternative Instrumental Variables

This table presents second-stage results from 2SLS regressions as robustness checks, using alternative instrumental variables to evaluate the impact of robot exposure on firms' innovation. In Panel A, we replace the instrument with the average automation index of five European countries plus Germany. Panel B uses the average automation index of nine European countries—Denmark, Finland, France, Germany, Italy, Norway, Spain, Sweden, and the United Kingdom—as the instrument. In Panel C, we address potential time-related confounding by regressing the original instrument on time-fixed effects and using the residuals as an alternative measure of robot exposure. The dependent variables include: *Tech Similarity*, *Semantic Similarity*, *AI Closeness*, *AI Distance*, *R&D/Total Assets*, *R&D/Sales*, and *Patent Value*. *Robot Exposure* quantifies the degree of automation faced by each firm in a given year, standardized to have a mean of zero and a standard deviation of one. Firm-level control variables include firm size (measured by total assets), firm age, leverage ratio, return on assets, and the proportion of tangible assets. The sample period spans 2004–2019. The regression data are at the firm-year level. All regressions incorporate firm-fixed and year-fixed effects. Firm-level clustered standard errors are reported in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Level	Firm-Year						
Period	t						
Model	IV						
Variable	Tech Similarity	Semantic Similarity	AI Closeness	AI Distance	R&D/Total Assets	R&D/ Sales	Patent Value
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Panel A: IV: 6 European Countries						
Robot Exposure	-0.300*	-0.436**	0.004**	-0.736**	0.181**	0.472**	-0.997**
	(0.180)	(0.215)	(0.002)	(0.333)	(0.078)	(0.204)	(0.475)
Observations	20,815	18,238	20,815	20,815	20,815	20,815	20,815
F Statistics	24.000	18.740	24.000	24.000	24.000	24.000	24.000
P-value (Underidentification)	0.000	0.000	0.000	0.000	0.000	0.000	0.000
P-value (Endogeneity)	0.075	0.005	0.002	0.002	0.001	0.001	0.005
	Panel B: IV: 9 European Countries						
Robot Exposure	-0.366*	-0.517**	0.004*	-0.704**	0.149*	0.473**	-0.798*
	(0.205)	(0.261)	(0.002)	(0.359)	(0.081)	(0.221)	(0.482)
Observations	20,815	18,238	20,815	20,815	20,815	20,815	20,815
F Statistics	20.310	15.390	20.310	20.310	20.310	20.310	20.310
P-value (Underidentification)	0.001	0.001	0.001	0.001	0.001	0.001	0.001
P-value (Endogeneity)	0.042	0.002	0.011	0.005	0.031	0.001	0.035
	Panel C: IV: Adjusting for Time-Related Confounds						
Robot Exposure	-0.300*	-0.436**	0.004**	-0.736**	0.181**	0.472**	-0.997**
	(0.180)	(0.215)	(0.002)	(0.333)	(0.078)	(0.204)	(0.475)
Observations	20,815	18,238	20,815	20,815	20,815	20,815	20,815
F Statistics	24.000	18.740	24.000	24.000	24.000	24.000	24.000
P-value (Underidentification)	0.000	0.000	0.000	0.000	0.000	0.000	0.000
P-value (Endogeneity)	0.075	0.005	0.002	0.001	0.003	0.001	0.004
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table A15: Robustness Checks: Shift-share IV

This table presents second-stage results using the shift-share instrumental variable approach following Borusyak, Hull, and Jaravel (2022), with coefficients estimated through equivalent industry-level regressions to obtain exposure-robust standard errors. Panel A reports falsification test results based on industry-level and firm-level characteristics to assess the exogeneity of the European robot adoption shock. Panel A(a) includes industry-level control variables, including the share of high-tech equipment in total investment (*High-tech/Invest*), the ratio of input to value added (*Input/Value-added*), the log of real wages (ln(*Wage*)), wage per worker (*Wage/Worker*), the log of employment (ln(*Worker*)), and the ratio of workers to value added (*Worker/Value-added*). Panel A(b) incorporates firm-level control variables, including standard firm characteristics and a dummy variable indicating whether the firm has adopted an enterprise system (*ES*). Panel A(c) focuses on pre-trend variables measured over the period 1988–2003 to further validate the plausibility of the exclusion restriction. Panel B presents robustness check results using the shift-share IV to examine the impact of robot exposure on firms' innovation outcomes. The dependent variables include: *Tech Similarity*, *Semantic Similarity*, *AI Closeness*, *AI Distance*, *R&D/Total Assets*, *R&D/Sales*, and *Patent Value*. *Robot Exposure* measures the degree of automation faced by each firm in a given year, standardized to have a mean of zero and a standard deviation of one. The sample period covers 2004–2019. Firm-level clustered standard errors are reported in parentheses. Coefficients marked with *, **, and *** are significant at the 10%, 5%, and 1% levels, respectively.

Model	IV						
Period	t						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Panel A: Falsification Tests						
	(a) Industry-level Balance Tests						
Variable	High-tech/ Invest	Input/Value-added	Ln(Wage)	Wage/ Worker	Ln(Worker)	Worker/Value-added	
Robot Exposure	1.125	10.244	-3.778	126.877	-6.400	-29.470	
	(1.026)	(7.770)	(6.987)	(155.594)	(8.060)	(33.254)	
No. of Industry-Periods	208	208	208	208	208	208	
	(b) Firm-level Balance Tests						
Variable	Ln(Assets)	ROA	Leverage	Tangibility	Ln(Age)	ES	
Robot Exposure	-203.436	43.472	-23.099	7.516	-46.549	-24.738	
	(3,286.148)	(698.427)	(375.335)	(120.555)	(755.266)	(402.401)	
No. of Industry-Periods	208	208	208	208	208	208	
	(c) Firm-level Pre-trends Tests						
Variable Period	1988-2003						
Variable	Tech Similarity	Semantic Similarity	AI Closeness	AI Distance	R&D/Total Assets	R&D/ Sales	Patent Value
Robot Exposure (2004-2019)	0.185	-0.245	-0.009	0.418	0.046	0.089	-0.225
	(0.614)	(1.011)	(0.040)	(1.608)	(0.215)	(0.374)	(1.507)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
No. of Industry-Periods	60	60	60	60	60	60	60
	Panel B: Shift-share IV Results						
Variable	Tech Similarity	Semantic Similarity	AI Closeness	AI Distance	R&D/Total Assets	R&D/ Sales	Patent Value
Robot Exposure	-1.776	-1.166	0.049**	0.345	0.920*	0.090**	-22.698
	(2.783)	(1.867)	(0.021)	(0.246)	(0.490)	(0.041)	(35.917)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
No. of Industry-Periods	208	208	208	208	208	208	208

Table A16: Robustness of Text-Based Robot Exposure Measures

This table reports robustness checks related to the measurement of robot exposure and additional controls. Panels A–E use alternative measures of robot exposure, while Panel F adds a control for the readability of 10-K filings. In Panel A, we apply a two-stage regression procedure to refine and reweight keywords, with the resulting measure denoted as *Robot Exposure A*. In Panel B, we regress robot exposure on time-fixed effects and use the residuals from this regression as a cleaned measure, denoted as *Robot Exposure B*. In Panel C, we regress robot exposure on industry-fixed effects and use the residuals from this regression as a within-industry cleaned measure, denoted as *Robot Exposure C*. In Panel D, we construct a rolling weighted average of robot exposure from year $t-4$ to year t, applying weights of 0.2, 0.4, 0.6, 0.8, and 1, respectively. This measure is denoted as *Robot Exposure D*. In Panel E, we use the ratio of robots to employment as an alternative proxy, denoted as *Robot Exposure E*. The dependent variables include: *Tech Similarity*, *Semantic Similarity*, *AI Closeness*, *AI Distance*, *R&D/Total Assets*, *R&D/Sales*, *Patent Value*, and *Patent Count*. *Robot Exposure* quantifies the degree of automation faced by each firm in a given year, standardized to have a mean of zero and a standard deviation of one. Columns (1)–(7) present second-stage results from 2SLS regressions, while Column (8) reports estimates from a Poisson regression. Firm-level control variables include firm size (measured by total assets), firm age, leverage ratio, return on assets, and the proportion of tangible assets. The sample period spans 2004–2019. The regression data are at the firm-year level. All regressions incorporate firm-fixed and year-fixed effects. Firm-level clustered standard errors are reported in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Level	Firm-Year							
Period	t							
Model	IV							Poisson
Variable	Tech Similarity	Semantic Similarity	AI Closeness	AI Distance	R&D/Total Assets	R&D/ Sales	Patent Value	Patent Count
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Panel A: Robot Exposure: Refining and Reweighting Keywords							
Robot Exposure A	-0.697*	-0.835**	0.010**	-1.709**	0.421**	1.096**	-2.314**	-0.005
	(0.422)	(0.378)	(0.004)	(0.744)	(0.189)	(0.448)	(1.063)	(0.017)
Observations	20,815	18,238	20,815	20,815	20,815	20,815	20,815	16,468
F Statistics	14.130	13.730	14.130	14.130	14.130	14.130	14.130	
P-value (Underidentification)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	
P-value (Endogeneity)	0.057	0.004	0.001	0.002	0.002	0.000	0.005	
	Panel B: Robot Exposure: Adjusting for Time-Related Confounds							
Robot Exposure B	-0.198*	-0.287**	0.003**	-0.485**	0.119**	0.311**	-0.656**	-0.056**
	(0.118)	(0.141)	(0.001)	(0.219)	(0.052)	(0.134)	(0.313)	(0.023)
Observations	20,815	18,238	20,815	20,815	20,815	20,815	20,815	16,468
F Statistics	24.000	18.740	24.000	24.000	24.000	24.000	24.000	
P-value (Underidentification)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	
P-value (Endogeneity)	0.075	0.005	0.002	0.001	0.003	0.000	0.004	
	Panel C: Robot Exposure: Adjusting for Industry-Related Confounds							
Robot Exposure C	-0.300*	-0.436**	0.004**	-0.736**	0.181**	0.472**	-0.997**	-0.084**
	(0.180)	(0.215)	(0.002)	(0.333)	(0.078)	(0.204)	(0.475)	(0.034)
Observations	20,815	18,238	20,815	20,815	20,815	20,815	20,815	16,468
F Statistics	24.000	18.740	24.000	24.000	24.000	24.000	24.000	
P-value (Underidentification)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	
P-value (Endogeneity)	0.075	0.005	0.002	0.001	0.003	0.000	0.004	
	Panel D: Robot Exposure: Five-Year Rolling Weighted Average							
Robot Exposure D	-0.148*	-0.248**	0.002**	-0.363**	0.089**	0.233***	-0.492**	-0.040
	(0.081)	(0.109)	(0.001)	(0.148)	(0.035)	(0.089)	(0.216)	(0.033)
Observations	20,815	18,238	20,815	20,815	20,815	20,815	20,815	16,468
F Statistics	67.290	39.720	67.290	67.290	67.290	67.290	67.290	
P-value (Underidentification)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	
P-value (Endogeneity)	0.093	0.006	0.002	0.001	0.003	0.000	0.003	

(Continued)

Level	Firm-Year							
Period	t							
Model	IV							Poisson
Variable	Tech Similarity	Semantic Similarity	AI Closeness	AI Distance	R&D/Total Assets	R&D/ Sales	Patent Value	Patent Count
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Panel E: Robot Exposure: The Ratio of Robots to Employment							
Robot Exposure E	-0.018*	-0.027***	0.000***	-0.043***	0.011***	0.028***	-0.059***	0.015
	(0.010)	(0.010)	(0.000)	(0.015)	(0.004)	(0.009)	(0.023)	(0.049)
Observations	20,815	18,238	20,815	20,815	20,815	20,815	20,815	16,468
F Statistics	105.900	105.900	105.900	80.180	105.900	105.900	105.900	
P-value (Underidentification)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	
P-value (Endogeneity)	0.738	0.199	0.107	0.000	0.107	0.000	0.010	
	Panel F: Controlling for Readability of 10-K Filings							
Robot Exposure	-0.300*	-0.435**	0.004**	-0.742**	0.181**	0.472**	-0.985**	-0.085**
	(0.179)	(0.214)	(0.002)	(0.332)	(0.078)	(0.203)	(0.471)	(0.035)
Readability	-0.133	0.519	0.007	0.963	0.247	-1.221	0.126	-1.476
	(0.987)	(0.942)	(0.011)	(1.923)	(0.376)	(1.119)	(2.680)	(6.816)
Observations	20,750	18,174	20,750	20,750	20,750	20,750	20,750	16,408
F Statistics	24.32	18.97	24.32	24.32	24.32	24.32	24.32	
P-value (Underidentification)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	
P-value (Endogeneity)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table A17: Robustness Checks: The Impact of the *Alice* Supreme Court Decision on AI Patent Applications

This table presents robustness checks examining the impact of the U.S. Supreme Court's *Alice* decision on AI-related patent applications. Panel A investigates how the *Alice* ruling affects AI patent filings using two outcome variables. The first, $Dummy(Grant)$, is a binary variable equal to 1 if a patent is ultimately granted, and 0 otherwise. The second, $Ln(Grant\ Lag)$, is the natural logarithm of the number of days between patent application and approval, plus one. We include the interaction term $Dummy(AI\ Patent) \times Dummy(After\ 2014)$ to capture differential changes in patent outcomes between AI and non-AI patents before and after the *Alice* decision (i.e., post-2014). Here, $Dummy(AI\ Patent)$ equals 1 for AI patents and 0 otherwise, and $Dummy(After\ 2014)$ equals 1 for patents filed in 2014 or later, and 0 otherwise. Columns (1)–(3) report OLS estimates, while Column (4) presents Poisson regression results. Panel B includes $Dummy(AI\ Patent\ \&\ After\ 2014)$ as a control variable, which equals 1 for AI patents filed in 2014 or later, and 0 otherwise, and reports the second-stage results from 2SLS regressions. All regressions include firm-fixed and year-fixed effects, and control for firm-level characteristics: firm size (total assets), firm age, leverage ratio, return on assets, and the proportion of tangible assets. The sample covers the period from 2004 to 2019, and the unit of observation is the patent-firm-year. Firm-level clustered standard errors are reported in parentheses. Coefficients marked with *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Level	Patent-Firm-Year			
Period	t			
	(1)	(2)	(3)	(4)
	Panel A: The Effects of Alice on AI Patent Approvals			
Model	OLS			Poisson
Variable	Dummy(Grant)		Ln(Grant Lag)	Grant Lag
Dummy (AI Patent) × Dummy (After 2014)	-0.099	-0.032	0.018**	0.038*
	(0.000)	(0.000)	(0.009)	(0.023)
Dummy (AI Patent)	0.076	0.050	0.027***	0.024
	(0.000)	(0.000)	(0.007)	(0.023)
Dummy (After 2014)	0.084			
	(0.000)			
Controls	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	Yes	Yes
Observations	3,340,083	3,216,325	659,190	659,185
	Panel B: Controlling for the Impact of Alice Decision			
Model	IV			
Variable	AI Closeness	AI Distance	Generality	Originality
Robot Exposure	0.082***	-1.027***	0.042***	0.041***
	(0.003)	(0.123)	(0.005)	(0.004)
Dummy (AI Patent & After 2014)	0.690***	-4.941***	0.002	0.002
	(0.004)	(0.076)	(0.006)	(0.005)
Controls	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	607,772	607,772	365,009	607,772
F Statistics	14.540	14.540	17.290	15.010
P-value (Underidentification)	0.000	0.000	0.000	0.000
P-value (Endogeneity)	0.000	0.000	0.000	0.000

Table A18: Robustness Checks: Alternative Approaches to Sample Construction

This table reports robustness checks using alternative sample construction approaches. Specifically, Panel A drops firms in the Internet industry, Panel B drops firms that have implemented enterprise systems, Panel C drops firms whose overseas sales exceed 50% of total sales, and Panel D drops the top 1% of firms with the highest robot exposure within each industry-year. The dependent variables include: *Tech Similarity*, *Semantic Similarity*, *AI Closeness*, *AI Distance*, *R&D/Total Assets*, *R&D/Sales*, *Patent Value*, and *Patent Count*. *Robot Exposure* quantifies the degree of automation faced by each firm in a given year, standardized to have a mean of zero and a standard deviation of one. Columns (1)–(7) present second-stage results from 2SLS regressions, while Column (8) reports estimates from a Poisson regression. All regressions include firm-fixed and year-fixed effects, and control for firm-level characteristics: firm size (total assets), firm age, leverage ratio, return on assets, and the proportion of tangible assets. The sample covers the period from 2004 to 2019. The unit of observation is the firm-year. Firm-level clustered standard errors are reported in parentheses. Coefficients marked with *, **, and *** are statistically significant at the 10%, 5%, and 1% levels, respectively.

Level	Firm-Year							
Period	t							
Model	IV							Poisson
Variable	Tech Similarity	Semantic Similarity	AI Closeness	AI Distance	R&D/Total Assets	R&D/ Sales	Patent Value	Patent Count
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Panel A: Drop Firms in the IT Industry							
Robot Exposure	-0.295	-0.425*	0.004**	-0.653**	0.189**	0.501**	-0.997**	-0.083**
	(0.184)	(0.220)	(0.002)	(0.326)	(0.082)	(0.215)	(0.490)	(0.039)
Observations	19,290	16,833	19,290	19,290	19,290	19,290	19,290	15,865
F Statistics	21.950	16.810	21.950	21.950	21.950	21.950	21.950	
P-value (Underidentification)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	
P-value (Endogeneity)	0.090	0.008	0.002	0.007	0.003	0.000	0.006	
	Panel B: Drop Firms with ES Systems							
Robot Exposure	-0.298	-0.501*	0.001	-1.004**	0.189**	0.555**	-1.093*	-0.095***
	(0.216)	(0.281)	(0.001)	(0.464)	(0.094)	(0.267)	(0.571)	(0.035)
Observations	18,654	16,228	18,654	18,654	18,654	18,654	18,654	15,354
F Statistics	18.050	13.420	18.050	18.050	18.050	18.050	18.050	
P-value (Underidentification)	0.002	0.006	0.002	0.002	0.002	0.002	0.002	
P-value (Endogeneity)	0.137	0.008	0.002	0.000	0.010	0.001	0.006	
	Panel C: Drop Firms with Overseas Sales $>$ 50%							
Robot Exposure	-0.300*	-0.436**	0.004**	-0.736**	0.181**	0.472**	-0.997**	-0.083**
	(0.180)	(0.215)	(0.002)	(0.333)	(0.078)	(0.204)	(0.475)	(0.039)
Observations	20,815	18,238	20,815	20,815	20,815	20,815	20,815	18,654
F Statistics	24.000	24.000	24.000	24.000	24.000	24.000	24.000	
P-value (Underidentification)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	
P-value (Endogeneity)	0.076	0.005	0.002	0.002	0.003	0.001	0.005	
	Panel D: Drop Firms with Top 1% Robot Exposure within each Industry-year							
Robot Exposure	-0.298	-0.434**	0.004**	-0.747**	0.185**	0.480**	-1.015**	-0.085**
	(0.183)	(0.219)	(0.002)	(0.341)	(0.080)	(0.209)	(0.486)	(0.034)
Observations	20,805	18,228	20,805	20,805	20,805	20,805	20,805	16,468
F Statistics	23.250	18.050	23.250	23.250	23.250	23.250	23.250	
P-value (Underidentification)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	
P-value (Endogeneity)	0.082	0.006	0.002	0.002	0.003	0.001	0.005	
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Appendix C: Derivations and Proofs

The first order condition for the instantaneous profit implies that

$$q_M^{\beta}(f)k_M^{-\beta}(s \mid f) = \frac{1}{\chi_M} \quad \forall s \in [0, \theta_f] \tag{28}$$

$$q_L^{\beta}(f)k_L^{-\beta}(s \mid f) = \frac{w}{\bar{q}\chi_L} \quad \forall s \in [\theta_f, 1] \tag{29}$$

Thus, the demand for

$$k_M(s \mid f) = q_M(f)\chi_M^{1/\beta}\ ,\ k_L(s \mid f) = q_L(f)\left(\frac{\bar{q}\chi_L}{w}\right)^{1/\beta},$$

and the demand for labor is

$$l(s \mid f) = \frac{q_L(f)}{\chi_L \bar{q}} \cdot \left(\frac{\bar{q}\chi_L}{w}\right)^{1/\beta}$$

Market clear for labor implies

$$w = \bar{q}_L \chi_L^{1-\beta}. \tag{30}$$

Therefore, firm f's profit conditional on the technology $q_M(f)$ and $q_L(f)$ takes the form of

$$\Pi_f = \alpha_M \theta_f q_n(f) + \alpha_L(1-\theta_f)q_L(f), \tag{31}$$

where

$$\alpha_M = \frac{\beta}{1-\beta}\chi_M^{\frac{1-\beta}{\beta}} \text{ and } \alpha_L = \frac{\beta}{1-\beta}\chi_L^{1-\beta}.$$

To determine equilibrium wage, we note that there is one unit of labor

$$1 \equiv \int_0^1 \left[\int_{\theta_f}^1 l(s \mid f)ds\right] df = \int_0^1 \frac{q_L(f)}{\chi_L \bar{q}}(1-\theta_f)\left(\frac{\bar{q}_L\chi_L}{w}\right)^{1/\beta} df \tag{32}$$

The profit function

$$\begin{aligned}
\Pi_f &= \frac{\beta}{1-\beta}\theta_f \chi_M^{\frac{1-\beta}{\beta}} q_M(f) + \frac{\beta}{1-\beta}(1-\theta_f)\left(\frac{\bar{q}_L\chi_L}{w}^{\frac{1-\beta}{\beta}}\right) q_L(f) \\
&= \frac{\beta}{1-\beta}\chi_M^{\frac{1-\beta}{\beta}}\theta_f q_M(f) + \frac{\beta}{1-\beta}\chi_L^{1-\beta}(1-\theta_f)\cdot q_L(f) \\
&= \alpha_M \theta_f q_M(f) + \alpha_L(1-\theta_f)q_L(f)
\end{aligned} \tag{33}$$

with

$$\alpha_M = \frac{\beta}{1-\beta}\chi_M^{\frac{1-\beta}{\beta}}, \quad \alpha_L = \frac{\beta}{1-\beta}\chi_L^{1-\beta}.$$

Here, we have a setup with

$$\Pi_f = \alpha_M \theta_f q_M(f) + \alpha_L (1-\theta_f) q_L(f) \tag{34}$$

We obtain the solution

$$\begin{aligned} &\gamma A_M q_M = \max_{z_M} \left\{ \alpha_M \theta_f q_M - C_M(z_M) q_M + z_M A_M \lambda_M q_M \right\} \\ &\Rightarrow \\ &\gamma A_M = \max_{z_M} \left\{ \alpha_M \theta_f - C_M(z_M) + z_M A_M \lambda_M \right\} \end{aligned} \tag{35}$$

Similarly,

$$\gamma A_L = \max_{z_L} \left\{ \alpha_L (1-\theta_f) - C_L(z_L) + z_L A_L \lambda_L \right\} \tag{36}$$

The first-order conditions (F.O.C.) are:

$$\begin{cases} C_M'(z_M^*) = A_M \lambda_M, \\ C_L'(z_L^*) = \lambda_L A_L. \end{cases} \tag{37}$$

A_M, A_L are determined.
Now, consider a firm's robot exposure. For a firm f, robot exposure is defined as:

$$\text{Robot exposure} f = \frac{\int_0^{\theta_f} k_M(s \mid f) ds}{\int_{\theta_f}^1 k_L(s \mid f) ds} = \frac{\chi_M^{1/\beta}}{\chi_L} \frac{q_M(f)}{q_L(f)} \frac{\theta_f}{1-\theta_f} \tag{38}$$

We examine how automation degree θ_f affects innovation strategy

$$\frac{\partial z_M^*}{\partial \theta_f} \quad \frac{\partial z_L^*}{\partial \theta_f} \tag{39}$$

Using the envelope theorem:

$$\begin{aligned} \gamma \cdot \frac{\partial A_M}{\partial \theta_f} &= \alpha_M, \\ \gamma \cdot \frac{\partial A_L}{\partial \theta_f} &= -\alpha_L, \\ \frac{\partial z_M^*}{\partial \theta_f} &= \frac{\lambda_M}{C_M''(z_M)} \cdot \frac{\partial A_M}{\partial \theta_f} = \frac{\lambda_M}{C_M''(z_M)} \cdot \frac{\alpha_M}{\gamma} > 0. \end{aligned} \tag{40}$$

$\lambda_M, \alpha_M, C_M''(z_M)$ represent heterogeneity analysis.
We also have:

$$\alpha_M = \frac{\beta}{1-\beta} \cdot \chi_M^{\frac{1-\beta}{\beta}}, \quad \alpha_L = \frac{\beta}{1-\beta} \chi_L^{1-\beta}$$

and $C''(z_M)$ measures the cost increase.
Finally,

$$\frac{\partial z_L^*}{\partial \theta_f} = \frac{\lambda_L}{C_L''(z_L)} \cdot \frac{\partial A_L}{\partial \theta_f} = -\frac{\lambda_L}{C_L''(z_L)} \cdot \frac{\alpha_L}{\gamma} < 0. \tag{41}$$