Automation-Induced Innovation Shift*

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

May 2024

Abstract

We study the impact of exposure to automation on corporate innovation, which informs how innovation begets innovation. We document that firms with high robot exposure witness a decline in technology similarity over time and significantly shift innovative activities toward AI, which automation intuitively complements. The shift is more pronounced for firms with greater AI-related research experience or more data generation. Furthermore, AI patents are more costly than non-AI patents in various dimensions, including team size, researchers' labor input and inventors' originality. Consequently, firms with high automation experience a significant rise in R&D expenditure but a decline in the number of new patents in subsequent years. Against the backdrop of rising automation, this explains the puzzling observation in the literature that at the aggregate level, firms seem to become less innovative in recent years despite greater R&D expenditures. Finally, we present a simple dynamic equilibrium model to rationalize such innovation shifts.

JEL Classification: G30, G40
Keywords: AI, Corporate Innovation, Language Models, Technology, Robot

*We thank Anastassia Fedyk, Yueran Ma, Xin Chang (Simba), Hong Zhang, Haotian Xiang, Siguang Li, and Pengfei Han, Gerard Hoberg for their insightful comments on our paper and participants at the 2025 ASSA annual meeting (Econometric Society), Five-star conference in finance, AsianFA, and seminar participants in GSM Peking University, HKUST (GZ), and Tsinghua University. Lu acknowledges financial support from the Humanities and Social Sciences Major Project of the Ministry of Education (No. 22JJD790047), the "Double-High" Program for the 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). Zhu acknowledges financial support from Tsinghua University Initiative Scientific Research Program (No.2022Z04W02016) and Tsinghua University School of Economics and Management Research Grant (No.2022051002). All errors are our own.

†Cornell SC Johnson College of Business (Johnson) & NBER, will.cong@cornell.edu

‡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

1 Introduction

The transformative wave of automation, heralded as the Fourth Industrial Revolution by both media and industry leaders, has shaped economic developments around the globe and become a focal point of social discourse concerning its implications for labor markets, including effects on labor share, job creation, wages, and economic inequality.[1] However, we understand little how automation and robotics as a general purpose technology influences the core drivers of economic growth: the innovative activities of firms [Romer, 1986, Aghion and Howitt, 1992, Schumpeter and Backhaus, 2003]. This article aims to bridge the gap by investigating for the first time how a firm's exposure to automation and robotics dynamically incluences the direction, composition, and intensity of its innovation.

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. Such a standardized metric for robotic technology adoption and integration at the firm level overcomes the limitations of industry-level proxies used in prior studies [e.g., Zhang, 2019, Acemoglu and Restrepo, 2020]. Using an instrumental variable induced by the demographic change of 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 activities and a sharp rise in R&D expenditures. Upon closer scrutiny, the short-term reduction in innovation outputs is offset by a greater originality and generality of the innovation, underscoring the substantial resource reallocation required for AI-driven innovation and implying an innovation J-curve. Our paper therefore offers initial insights on how innovation begets innovation.

Specifically, we first examine broader trends in firms' technological shifts using two complementary metrics: (i) *Technology similarity* captures shifts in a firm's patent portfolio composition by representing each portfolio as a vector, where dimensions correspond to the firm's patent activity across various classifications in the USPTO. To track the evolution of a firm's

[1] see https://www.forbes.com/sites/benjaminlaker/2022/11/17/what-leaders-should-know-about-emerging-technologies/?sh=16da588262d1, and Acemoglu et al. [2018], Acemoglu and Restrepo [2020], Koch et al. [2021].

technological focus, we construct annual patent distribution vectors over a five-year horizon and apply cosine similarity to assess changes over time. (ii) *Semantic similarity* leverages PatentBERT—a variant of the BERT large language model specifically fine-tuned for the patent domain—to generate semantic representations of patent texts. This approach enables the detection of 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 exposure to robotics experience a significant 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 instrument introduced by Borusyak et al. [2022].

Next, we focus on Artificial Intelligence (AI)—a key technological domain toward which firms increasingly reorient their innovations in response to rising robot exposure. This shift is driven by two key factors: AI enhances robotic functionality by improving operational capabilities, while simultaneously benefiting from the vast data generated by robots, fostering a self-reinforcing cycle of AI-driven advancements. To further substantiate this trend, we visualize the relative shifts in innovation direction across firms. The results indicate that firms with high robot exposure increasingly patent in fields closely linked to AI, providing preliminary evidence supporting our hypothesis of AI as a primary area of technological adaptation.

To explore AI-focused innovation, we employ text analysis to empirically identify AI-related patents, isolating those whose abstracts contain AI-specific keywords as Cockburn et al. [2018]. This approach overcomes the limitations of cooperative patent classification (CPC), which does not explicitly distinguish between AI and non-AI patents. Beyond a binary AI versus non-AI classification, we develop two novel metrics: AI distance and AI closeness. 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. This approach leverages a patent-to-patent citation network, where: AI distance captures the citation network distance between a given patent and the representative AI patent and AI closeness measures the

similarity in citation patterns between a given patent and the AI archetype, providing a more granular assessment of AI relevance (detailed in Section 3.3). 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 pattern: greater robot exposure is associated with a significant reduction in patent distance from AI innovations, alongside a corresponding increase in patent closeness to AI-related technologies. Additionally, firms with high robot exposure are more likely to produce a greater number of AI-related patents in subsequent years, indicating a strategic shift toward AI-driven innovation. To examine the underlying mechanism driving this transition, we investigate the role of enhanced data access. Our analysis shows that firms with heightened robot exposure experience a notable surge in data generation activities, as evidenced by our text analysis. This empirical evidence suggests that firms strategically reorient their innovation efforts toward AI as a direct response to the integration of robotics, leveraging data-driven insights to accelerate AI advancements.

For robustness as well as more granular analyses, we exploit the AI dataset by Giczy et al. [2022] to understand which AI subclasses firms increasingly focus on in response to heightened robot exposure. This dataset constructed using a machine learning model to classify U.S. patents categorizes AI innovations into eight subclasses: machine learning, evolutionary algorithms, natural language processing, speech processing, vision processing, knowledge representation, planning and control algorithms, and hardware. We find that firms with high robot exposure experience a significant innovation increase across all eight AI subclasses, with the strongest effects in knowledge representation, planning and control algorithms, and hardware, which are crucial for robotics integration.

In addition to the direction of innovation, we further explore how automation exposure influences overall innovative activities, including both their inputs and outputs, which is unclear a priori. Concerning innovation outputs, we identify an innovation "J-curve," characterized by an initial decline followed by a subsequent increase in the citation-weighted number of patents. Specifically, as robot exposure intensifies, the cumulative number of patents undergoes a significant decline in the first two years, though the rate of decline gradually diminishes over time. By the fifth year, the coefficient on robot exposure transitions from negative to positive, culmi-

nating in a substantial increase in the number of patents by the sixth year. This pattern suggests that firms initially experience a contraction in patent output 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 significant uptick in R&D investments as well as a reduction in patent outputs in the years following increased automation exposure. 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 19.2% and 62.6% increases in the R&D-to-total asset and R&D-to-sale ratios, respectively. These shifts are statistically significant and bear substantial economic implications, especially when contrasted with the mean R&D-to-total asset and R&D-to-sale ratios of 46.2% and 48.2%. The findings are also consistent with Babina et al. [2024], who demonstrate that AI investments, as slow-moving processes, stimulate firms' long-term innovation output through increased R&D expenditures rather than generating immediate effects.

This puzzling gap between R&D spending and patenting may be attributed to more foundational research with greater originality and generality, the high transition costs of specialized human or fixed capital investments required for new technology fields or 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, are costly but enables breakthroughs in foundational science and technology Jumper et al. [2021]. However, 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 less, albeit higher-quality, patents. Adopting the methodology proposed by Hall et al. [2001], we indeed find that AI patents exhibit significantly higher levels of both originality and generality.

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.

This hypothesis is empirically validated as our analysis reveals that firms with extensive AI research history exhibit more significant innovation shifts, as evidenced by increased closeness to and decreased distance from AI-related patents. Our exploration delves into the nuanced cost dynamics of AI-related innovation compared to non-AI fields, employing a series of metrics designed to gauge the marginal costs of innovation as proposed by Jones [2009] who develop a theory to explore the implication of technology-specific human capital on technology progress. Our analysis reveals that AI innovation, in contrast to their non-AI counterparts, demands more substantial team collaboration, a greater volume of growth inventors' labor input, and exhibits higher levels of inventor originality, further corroborating the idea that automation exposure spurs firms to focus more on fundamental and bigger innovations, which requires more inputs and takes longer to generate outputs.

Is such an innovation shift persistent in the longer run? First, we find that the shift toward AI-driven innovation exhibits lasting impacts, with AI closeness remaining significantly positive after ten years. The effect on innovation output also persists, as reflected in the citation-deflated patent number, which stays significantly negative at the 10% level after eight periods. In contrast, the influence of robot exposure on other variables is predominantly short-term, lasting no more than four periods. That said, our findings indicate that the magnitude of the effect diminishes over time, ultimately disappearing over the longer horizon. This supports the notion that the initial decline is driven by transitory adjustment costs. There is a notable increase in the number of patents in the sixth year, corroborating the innovation J-Curve we hypothesize.

We demonstrate the robustness of our core findings through an array of additional tests involving alternative constructions of key variables, controlling for Alice decisions and other technological shocks, varying sample specifications, testing results at different levels, performing placebo tests, etc. We also build a dynamic general equilibrium model to understand the underlying forces driving innovation shifts amidst increasing automation and to rationalize the aforementioned empirical observations. The model accommodates firms with varying degrees of automation exposure and allows for endogenous innovation choices. Firms operate a spectrum of tasks, with a portion being automated and the remainder executed 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.

Our study adds to studies concerning the dimensions of technological progress. Prior empirical inquiries have predominantly focused on how market size expansion influences the development of new products or the dynamics between input factors that complement or substitute specific production factors. For instance, Finkelstein [2004] highlight the stimulative impact of public policy on vaccine-related clinical trials. Similarly, Beraja et al. [2022] illustrate how government security procurement contracts in China have spurred AI development by providing government-collected data. Furthermore, research by Acemoglu and Linn [2004] and Costinot et al. [2019] show that demographic shifts can expand markets and drive innovation in pharmaceuticals. In addition, Acemoglu and Restrepo [2022] provided international evidence on how demographic changes can accelerate automation adoption, while Bloom et al. [2021] identified a shift in patent applications towards work-from-home technologies triggered by the COVID-19 pandemic. We contribute by (i) underscoring the novel role of technology adoption, specifically the recent adoption of automation, in catalyzing directed innovation shifts, and (ii) illuminating the innovation "J-Curve" entailing a reduction in patent output within incumbent fields, a phenomenon previously undocumented but related to the recent discussion of technological obsolescence [Ma, 2023].

Our investigation 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 an increasing demand for knowledge and learning for inventors to approach the innovation frontier over time. Bloom et al. [2020] employs industry-level case studies to similarly illustrate a significant rise in research efforts juxtaposed against a stark decline 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 endogenous innovation shift, while strategically beneficial in the long run, may manifest in the seeming reduction of research productivity observed across industries.

Finally, our study enriches the discourse on the determinants of corporate innovation, which

entails topics ranging from institutional and legal frameworks to market structures and individual corporate attributes (Ederer and Manso [2011], He and Tian [2018], Lin et al. [2021], Hegde et al. [2022]). Notably, significant emphasis has been placed on the influence of broad economic forces, including intellectual property regimes (Lerner [2009], Fang et al. [2017]), labor and bankruptcy laws (Acharya et al. [2014], Acharya and Subramanian [2009], Cerqueiro et al. [2017]), and the nuances of financial market evolution (Hsu et al. [2014]), alongside market characteristics (Bloom et al. [2013], Yung [2016], Autor and Salomons [2018]) and specific corporate traits (Malmendier and Tate [2005], Lerner [2000], He and Tian [2013], Bena and Li [2014], Sunder et al. [2017]) in shaping the innovation landscape.

While these studies provide invaluable insights about drivers of innovation, there remains a paucity of research on how technological breakthroughs directly influence firms' innovation strategies. Our work fills the gap by focusing on automation as a byproduct of technological progress and as a pivotal force reshaping industry-wide innovation patterns. Beyond conventional analysis centered on aggregate patent output, our novel metrics for tracing the evolution of patent distribution across technological domains enable us to explore how innovation shifts reallocate patent activity among different technological fields. Through a detailed examination of how firms recalibrate their innovation directions in response to the ascendancy of automation, we uncover nuanced insights into the dynamics of patent numbers and innovation focus, thereby adding a critical dimension to the understanding of how technological advancements directly impact corporate innovation trajectories.

The remainder of the paper is organized as follows. Section 2 describes the data and empirical strategy. Section 3 presents our baseline results of directed innovation shift. Section 4 presents the empirical findings on aggregate innovation, while Section 5 discusses robustness. Section 6 introduces the theoretical model, before Section 7 concludes.

2 Data, Variables, and Empirical Strategy

2.1 Data Description

Our data primarily come from the following databases:

(i) Patent-Related Data. Patent-related data are obtained from the United States Patent and Trademark Office (USPTO), which has provided comprehensive data on patents granted in the United States since 1976.[2] Our study utilizes patent applications and grants to measure firms' innovation output. Given the considerable time lag between a patent application and its eventual grant—often spanning several years—following Kogan et al. [2017], we use the application date to construct our annual patent count, as a more timely reflection of innovation activities. To match patent data with firm-level financial information, we employ the linkage methodology developed by Kogan et al. [2017].

(ii) Robotics Data. We source robotics-related data from two primary channels. At the firm level, we extract keywords used to assess robot exposure from the 10-K filings of U.S. publicly traded firms. At the industry level, we obtain data from the International Federation of Robotics (IFR), which provides detailed records on robot stocks compiled from major robot manufacturers worldwide in each year. The IFR dataset covers approximately 90 percent of industrial robots, including those deployed in the United States, Finland, Sweden, Denmark, Italy, and France.

(iii) Firm-Level Financial Data. Firm-level financial and operational data are obtained from the Compustat database. This dataset includes key variables such as R&D intensity measures (R&D/Total Assets, R&D/Sales), firm size (Ln(Assets)), leverage, firm age (Ln(Age)), profitability (ROA), and asset tangibility. Additionally, we use segment data from Compustat to extract firms' industry classification and revenue composition.

(iv) Industry-Level Data. We obtain industry-level data on wages, employment, and prices from the Bureau of Labor Statistics (BLS), while data on investments, inputs, and value-added come from the Bureau of Economic Analysis (BEA). To construct an instrumental variable, we leverage European employment statistics from the European Union–level analysis of capital, labor, energy, materials, and service inputs (EU KLEMS) Growth and Productivity Accounts, which provide detailed information on capital, labor, energy, materials, and service inputs.

[2]Detailed information can be accessed at USPTO Patent Data.

Our data span the period from 2004 to 2019. We select 2004 as the starting year because IFR data indicate that, prior to this point, industrial robot installations across U.S. industries were recorded as zero,[3] effectively marking 2004 as the inaugural year of industrial robot adoption in the United States. The endpoint, 2019, corresponds to the latest available IFR data at the time of our analyses.

To maintain a consistent sample, we exclude firms from the finance and utility sectors, as well as non-domestic entities identified by their country of incorporation, due to the unique regulatory and operational environments of these firms. Additionally, we restrict our analysis to firms that have applied for at least one patent during the sample period. We further impose a completeness requirement on core financial variables[4] to ensure data reliability.

Table 1 provides an overview of our sample composition, detailing the number of U.S. public firms included, the annual percentage of firms with at least one patent, and the total number of patents attributed to this sample. Notably, approximately half of the firms in our dataset receive at least one patent annually, underscoring a high level of innovative activity.

2.2 Automation and Robots

As data on industrial robot installations are available exclusively at the industry level, prior studies such as Zhang [2019] and Acemoglu and Restrepo [2020] estimate firm-level robot exposure by leveraging industry-level data from the IFR. Specifically, they calculate robot exposure as the number of robot installations per thousand workers, combined with Compustat segment data. Firm-level robot exposure is computed as:

$$\text{Robot Exposure (IFR)}_{i,t} = \sum_{j=1}^{m} \frac{\text{Sales}_{i,j,t}}{\text{Sales}_{i,t}} \times \frac{\text{Robots}_{j,t}}{\text{Workers}_{j,t}}, \tag{1}$$

where $\text{Sales}_{i,j,t}$ represents the annual sales of firm i in industry segment j in year t, and $\text{Sales}_{i,t}$ denotes total annual sales of firm i in year t. The term $\frac{\text{Robots}_{j,t}}{\text{Workers}_{j,t}}$ captures industry-level robot exposure in segment j for year t. However, given that many firms, as classified by Compustat,

[3] With the exception of a single unspecified industry that reported installations.

[4] Core financial variables refer to the control variables used later in the analysis, including firm size, leverage, firm age, profitability, and asset tangibility.

primarily operate within a single segment, this measure often reflects industry-level exposure rather than firm-specific interactions with robotic technologies, potentially limiting its precision in capturing firm-level technological adoption.

To address this limitation, we leverage textual analysis of 10-K filings, which provide detailed firm-level disclosures on business activities and technological advancements. Specifically, we construct a firm-level robot exposure measure by identifying mentions of robot-related keywords within 10-K filings from 2004 to 2019. The process involves two key steps. First, we compile a comprehensive keyword list based on prior literature [Zeira, 1998, Acemoglu and Restrepo, 2018, 2019, 2020, 2022, Berg et al., 2018, Graetz and Michaels, 2018, Caselli and Manning, 2019, Dixon et al., 2021, Koch et al., 2021, Guerreiro et al., 2022] and refine it using definitions from IFR reports.[5] Second, we exclude terms related to the sale of industrial robots, ensuring that our focus remains on adoption and integration. Table 2 presents the finalized keyword list.

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

$$\text{Robot Exposure}_{i,t} = \sum_{w} \frac{\text{Robot Word Count}_{w,i,t}}{\text{Total Word Count}_{i,t}}, \tag{2}$$

where $\text{Robot Word Count}_{w,i,t}$ represents the occurrences of keyword w for firm i in year t, and $\text{Total Word Count}_{i,t}$ is the total number of words in the corresponding report. This normalization enables consistent comparisons of robot exposure across firms. 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. Figure 1 illustrates the relationship between our estimated robot exposure—aggregated to industry-level annual averages—and the industry-level annual averages of robot usage per worker in the U.S.[6] The results

[5] For details, see IFR Industrial Robots.

[6] To compute industry-level robot usage per worker, we utilize IFR data on robot usage and employment data from the BLS. Due to inconsistencies in industry classification between the IFR and BLS datasets, we follow the approach of Acemoglu and Restrepo [2020] and manually align 15 industries, including Agriculture, Forestry, and Fishing; Mining and Quarrying; Electricity, Gas, and Water Supply; Construction; Education; Research and Development; and nine manufacturing industries such as 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. Additionally, since IFR reports robot applications for the U.S., Canada, and

indicate a strong correlation between estimated robot exposure and actual robot usage, supporting the validity of our measure as a proxy for real-world robot adoption.

Tables A1 and A2 provide the definitions and summary statistics for robot exposure and other key variables used in subsequent analyses.

To ensure the robustness of our findings, we consider several alternative measure of robot exposure, with results presented in Section 5.2. First, we apply a two-stage Lasso regression [Cong et al., 2019] to select keywords that significantly predict industry-level robot installations. Second, we replicate the measure used by Zhang [2019] and Acemoglu and Restrepo [2020], utilizing industry-level robot installations per worker combined with Compustat segment data. Third, we regress robot exposure on time fixed effects and use the residual as the primary explanatory variable to filter out common time-varying factors influencing automation trends. Fourth, we employ a rolling five-year weighted average of robot exposure to account for factors such as depreciation. All results remain consistent across these specifications.

Beyond robustness, an additional concern is whether robot exposure follows a lumpy investment pattern—where firms make a large, one-time investment in robotic technology and then cease further investment in subsequent years. Humlum [2022] finds evidence of lumpy adoption in Danish firms that rely heavily on robot imports.[7]

Figure 2 illustrates the distribution and evolution of robot exposure among U.S. firms, revealing an overall upward trend with moderate year-to-year fluctuations. To further investigate this pattern, we analyze incremental changes in robot exposure. Each year, we identify the top 10% of firms experiencing the largest increases in robot exposure and track their robot exposure increment in the following year. Figure 3 shows that firms experiencing substantial shifts in robot exposure in one year often continue to see significant increases or decreases in subsequent years. This pattern suggests that investment in automation in the U.S. is not lumpy.

Mexico jointly at the industry level before 2011, we follow Acemoglu and Restrepo [2020] in including data from these three countries, as they document that the U.S. accounts for over 90 percent of North America's industrial robot market share.

[7]Using data from 454 Danish firms that use industrial robots, Humlum [2022] documents that 70.6% of firms in the sample invested in robots within the same year. However, this finding has not been widely confirmed in the broader literature. One possible explanation is that in countries such as Denmark, which rely heavily on robot imports, adoption patterns are influenced by factors such as import availability, supply chain dynamics, government policies, and technological upgrade cycles, leading to irregular adoption trends.

2.3 Identification Strategy and Instrumental Variable

To address potential endogeneity in analyzing the impact of automation exposure on corporate innovation, we confront two primary identification challenges. First, firms with active innovation efforts, particularly in AI, may disproportionately mention robotics within their 10-K filings. Second, the adoption of robotics may be a response to underlying structural changes or part of a broader technological evolution, potentially confounding the estimated impact on innovation outcomes.

To establish causality, we employ an instrumental variable (IV) strategy leveraging the aging population to workforce ratio in select European countries that are at the forefront of robotics technology. This IV builds on the work of Acemoglu and Restrepo [2020], who document that robotic adoption trends in these European nations are predictive of subsequent trends in the United States, and are driven by the demographic structure change. Notably, countries such as Denmark, Finland, France, Italy, and Sweden have demonstrated technological superiority in robotics compared to the United States.

We construct the IV using robot usage data from the IFR and employment data from the EU KLEMS database. Industry classification discrepancies between the IFR and EU KLEMS necessitate manual reconciliation, aligning 15 industries (SIC15) as detailed in Tables A3, A4, and A5. The IV is defined as:

$$Robot\ Exposure\ IV_{i,t} = \sum_{j=1}^{m} \frac{Sales_{i,j,t}}{Sales_{i,t}} \times \frac{Robots_{j,t}^{EU}}{Workers_{j,t}^{EU}}, \tag{3}$$

where $Robots_{j,t}^{EU}$ and $Workers_{j,t}^{EU}$ denote the number of industrial robots and workers, respectively, in industry j during year t in the EU countries mentioned above. Driven by the population aging, these countries have consistently maintained a higher density of robots since the 1990s and continue to invest in robotic technology (Acemoglu and Restrepo [2020]). For robustness, our analysis also incorporates data from Germany, Norway, Spain, and the United Kingdom.

We conduct a series of tests to validate the exclusion restriction of our IV—that robot adoption in European countries influences U.S. firm innovation only through its effect on U.S. robot exposure. First, Acemoglu and Restrepo [2022] demonstrate that the aging ratio is a crucial de-

terminant of the adoption rates of robots and varies significantly between countries. They show that the higher levels of robot adoption in these countries are largely driven by their more rapid aging shifts compared to the United States. These demographic trends in Europe, which differ from those in the United States, shape the automation adoption pattern, providing an external benchmark for U.S. firms.

Second, we empirically rule out potential channels through which robot adoption in European countries could affect U.S. firm innovation, such as industry-level price changes and other industry-year-specific factors. Moreover, as justified by Acemoglu and Restrepo [2020], industry-level robot exposure in the European Union should remain exogenous to firm-level innovation decisions in the United States.

Third, we substantiate the absence of significant pre-trends. By analyzing the relationship between robot exposure and innovation activity before 2004, we find no notable quantitative or statistical correlations. This examination ensures that the observed effects post-2004 are not merely a continuation of existing trends. Additionally, we control for variations in the dependent variable from 1988 to 2003 to fortify the robustness of our empirical outcomes. Detailed results, including analyses of pre-trends and alternative IV constructions, such as the shift-share IV method suggested by Borusyak et al. [2022] and a time-trend adjustment for the IV, are delineated in Section 5.1.

We also conduct a suite of validity tests for the instrumental variable, including the weak identification test, underidentification test, and endogeneity test. The results are reported in the first column of Table 3. The F-statistic for the weak identification test is 24.00, comfortably exceeding the critical value of 10 as suggested by Stock and Yogo [2005], thus mitigating the concerns of a weak instrument. Furthermore, the p-values for both the underidentification and endogeneity tests are below the conventional significance threshold of 0.1. These results indicate that our instrumental variable is adequately informative for parameter identification.

3 Robot Exposure and Innovation Shift

Our primary objective is to assess the influence of robot exposure on the innovation trajectory within firms, specifically investigating the extent to which rising exposure to robotics

catalyzes shifts in innovation. To provide a clearer visualization of the relative changes in innovation directions between firms, we compare the relative frequencies of various n-grams in the patent abstracts of firms with high and low robot exposure. The comparison is visualized in a word cloud[8] (Figure 4), indicating a trend in the patents of high-robot-exposure firms toward software, electronics, machinery, and energy—fields closely intertwined with AI.[9] Specifically, in 2010, n-grams such as "automated banking", "graphical program", "interface object", and "transmitting devices" emerge as early indicators of potential AI integration. By 2015, the prevalence of terms like "turbine engine", "gas turbine", "drill bit", and "downhole tool" signifies a continued emphasis on mechanical and energy technologies, which increasingly incorporate AI-driven real-time monitoring, predictive maintenance, and automated control. By 2019, the presence of terms such as "transmission medium", "syntax elements", "electromagnetic wave", "network controller", and "communications described" reflects a deeper convergence of AI with communication and information-transmission technologies, where AI plays a critical role in signal optimization, real-time decision-making, and large-scale data processing. This trajectory highlights the expanding role of AI across diverse industrial sectors.

Building on this visualization, we center our analysis on AI as a key technological domain toward which firms increasingly redirect their innovation efforts in response to rising robot exposure. This shift is driven by two fundamental mechanisms: AI innovations enhance and optimize the development and functionality of robotic systems, while robotics systems, in turn, generate vast data streams that train AI models and fuel further AI advancements. This mutual reinforcement establishes a self-perpetuating cycle, accelerating AI progress and deepening its integration into automation-driven innovation strategies.

Our subsequent analysis unfolds in three stages. First, we examine broader trends in techno-

[8]Specifically, we begin by preprocessing the patent texts, which includes converting all text to lowercase, removing HTML entities, and eliminating non-alphabetic characters (e.g., numbers and punctuation). Next, we extract 2-grams (or higher-order n-grams) from the preprocessed text, filtering out common stop words to retain meaningful word combinations and uncover potential contextual relationships. We then conduct a frequency analysis of the extracted n-grams, calculating their occurrences within the dataset. To reduce noise, we apply smoothing techniques, such as Laplace smoothing, to adjust frequency estimates and account for phrases that do not appear in the observed data. Subsequently, we compute the frequency values for patents in the high and low robot exposure quantiles for each year, relative to the baseline year of 2004. Finally, we derive the log-transformed relative frequencies between the two groups and visualize the resulting trends.

[9]Table A6 presents the top 20 n-grams with the highest relative frequency, corresponding to the results shown in Figure 4.

logical shift among firms. Next, we focus on the impact of robot exposure on firms' propensity to pivot their innovation efforts toward AI-related fields. Finally, we explore one potential mechanism through which increased robot exposure drives a shift toward AI innovation.

3.1 General Innovation Shift

To quantitatively measure the evolution of firms' technological focus, we employ a vectorization approach to represent each firm's patent innovations, utilizing cosine similarity as a metric to measure technological congruence. This analysis is conducted through two distinct vectorization methodologies to calculate similarity, each offering a unique lens on the firm's innovation activities.

Technology Similarity. To accurately capture the breadth of technological innovation across industries, patents are classified into technology classes based on their content and application, as recognized by the USPTO. The USPTO delineates these classes into hierarchical levels: 30 classes at the two-digit level for broad categorization, 126 at the three-digit level for more specific groupings, and 644 at the four-digit level for detailed classification. Our analysis focuses on the 126 three-digit CPC classes to ensure a balance between granularity and comprehensiveness. Let N represent the total number of technology classes. For a given firm i in year t, its patent distribution across these classes is denoted as $T_{it} = (T_{i1,t}, T_{i2,t}, \ldots, T_{iN,t})'$, where $T_{ij,t}$ represents the proportion of patents held by firm i in class j relative to its total patent count in year t. This distribution is based on patent applications filed by the firm between years $t-4$ and t to smooth year-to-year fluctuations, following the methodology of Bloom et al. [2013], Yang and Zhu [2020]. If a patent is classified under multiple classes, it is equally divided across them.

To quantify changes in a firm's technological focus over time, we define technology similarity as follows:

$$Tech\ Similarity_{i,t} = \frac{(T'_{i,t} T_{i,t+5})}{\sqrt{(T'_{i,t} T_{i,t})(T'_{i,t+5} T_{i,t+5})}} \tag{4}$$

Following Bloom et al. [2013], Yang and Zhu [2020], we employ cosine similarity between a firm's technology distribution vectors at time t and $t+5$ to measure technology similarity.

The use of a five-year window mitigates short-term fluctuations and provides a robust indicator of the firm's technological focus over this period. At its core, technology similarity measures how closely a firm's patent distribution across technology classes aligns over time, explicitly comparing the current year with five years into the future. A higher similarity score suggests a firm's consistent focus, while a lower score indicates a significant shift in the firm's innovation efforts across different technology classes. This measure provides a detailed understanding of a firm's innovation strategy and serves as a tool for anticipating and guiding technological transitions within firms.

Semantic Similarity. An alternative approach to measuring similarity involves analyzing changes in the abstracts of patents, as abstracts encapsulate the core themes of a firm's innovations. A straightforward method is to employ a bag-of-words (BoW) approach to quantify the distribution of words within patent abstracts. However, the BoW model treats words in isolation and disregards semantic relationships, limiting its effectiveness in capturing the nuanced meaning of patent abstracts. To address these limitations, we incorporate a more advanced natural language processing (NLP) technique that better accounts for semantic structures and contextual relationships within the text.

Specifically, we utilize PatentBert (Bekamiri et al. [2021]), a pre-trained BERT model fine-tuned on a comprehensive dataset of patent documents. This model effectively captures semantic relationships and contextual dependencies within patent abstracts. Each patent document is encoded into a 768-dimensional vector, offering a dense representation of its semantic content. We further refine these representations by incorporating patent citations, thereby accounting for variations in patent quality. To quantify shifts in technological focus, we return to the previously defined technology similarity metric (Equation (4)) and construct a BERT-based similarity measure. In this framework, T_{it} represents the citation-weighted distribution of words in the abstracts of patents granted to firm i in year t. This measurement also provides nuanced insights into how a firm's patenting activity evolves in response to strategic shifts toward emerging technological domains, with a lower BERT similarity indicating a more pronounced departure from the firm's previous technological trajectory.

To empirically test the hypothesis that higher robot exposure prompts firms to redirect their

innovation efforts towards new technological arenas, we employ the following specification:

$$y_{i,t+j} = \alpha_i + \alpha_t + \beta \times Robot\ exposure_{i,t} + X'_{it}\gamma + \varepsilon_{i,t} \tag{5}$$

Here, $y_{i,t+j}$ is measured using the cosine similarity between the firm i's technology distribution vectors from year $t+j$ to $t+j+5$ and the BERT-based patent document embeddings over the same period. $Robot\ exposure_{i,t}$ is measured as discussed in Section 2.2 and is standardized to have a mean of zero and a variance of one. α_i and α_t represent firm and year fixed effects, respectively, enabling within-firm and within-time comparisons. Firm fixed effects control for unobserved, time-invariant heterogeneity across firms, while year fixed effects capture common temporal shocks or trends (e.g., macroeconomic conditions or technological shifts). This specification isolates the variation of interest by accounting for both firm-specific and time-specific influences. X_{it} denotes the set of firm-level control variables, including firm size (log of total assets), leverage (debt-to-total assets ratio), firm age (log-transformed), profitability (ROA), and asset tangibility (tangible assets as a proportion of total assets). And $\varepsilon_{i,t}$ represents the error term. Error terms are adjusted for heteroskedasticity and clustered at the firm level to ensure robustness in our regression analysis.

Our empirical investigation, as detailed in Table 3, yields significant results regarding the effect of robot exposure on innovation shift captured by technology similarity as well as BERT similarity. The second-stage results of the IV estimation in Columns (2)–(6) reveals a substantial shift in the technological distribution over the subsequent years, driven by a rising robot exposure. Specifically, a one-standard-deviation increase in robot exposure is associated with a decline in technology distribution similarity of 0.187 in the current year, 0.213 in the following year, and 0.263 in the third year. Similarly, in terms of semantic similarity, the same increase in robot exposure leads to a reduction of 0.460 in the current period, 0.722 in the subsequent period, and 0.812 in the third period. These effects are both statistically significant and substantial, particularly given that the standard deviations of technology similarity and BERT similarity are 0.157 and 0.349, respectively. This finding highlights the transformative impact of robot exposure, driving firms to diversify their technological portfolios. Future research will build on these results by incorporating additional similarity measures to further explore how

robotic technologies reshape innovation trajectories.

3.2 AI-Oriented Innovation Shift

Initial findings indicate that firms with greater exposure to robot exhibit a markable shift in their innovation strategies, transitioning from established to emerging technological domains. We next explore the specific fields toward which these firms gravitate, particularly whether they pivot toward AI-related innovations that complement robotic technologies. To further support our empirical inferences, we present a dynamic general equilibrium model in Section 6, incorporating endogenous decisions regarding R&D expenditures and innovation. This model demonstrates that firms with heightened engagement in robotics naturally intensify their innovation efforts in areas that align with and complement robotic applications. And this section introduces our initial approach to measuring AI patents, laying the foundation for deeper exploration of AI-focused innovation.

A key challenge in identifying AI patents is the lack of explicitly defined AI-specific categories within the CPC system used by the USPTO, partly due to AI's broad and rapidly evolving nature. While the CPC framework systematically assigns patents to one or more classes and subclasses—ranging from Human Necessities (A) to Electricity (H), with a dedicated category (Y) for emerging technologies—it does not fully capture the complexity and interdisciplinary nature of AI. As a result, mapping patents to AI innovations based solely on classification codes remains challenging. To overcome this challenge, we adopt a text analysis approach focused on patent abstracts, allowing us to bypass the limitations of CPC classification by identifying AI-related patents based on AI-specific terminology. Specifically, we implement a text analysis protocol that classifies a patent as AI-related if its abstract includes at least one of the predefined keywords proposed by Cockburn et al. [2018]. These keywords encompass three core AI domains—*Robotics, Symbolic Systems, and Neural Networks*—as detailed in Table 4. To validate the effectiveness of this keyword-based approach for identifying AI patents, we conduct a robustness check in Section 3.4 using the AI patent dataset compiled by the USPTO. To the best of our knowledge, this dataset is the only publicly available classification of AI patents, although its accuracy remains subject to certain limitations.

Our analysis in Figure 4 suggests a noticeable shift in patent activity toward AI-oriented innovations between 2004 and 2019, highlighting the increasing emphasis on AI technologies. To further examine the temporal dynamics of AI-related innovation, we analyze the evolution of the relative frequency of AI-related keywords over time. We begin by aggregating patent abstracts annually and computing the occurrence of each AI keyword within that year's patent pool. The frequency of a keyword w in a given year t is defined as:

$$Frequency_{w,t} = \frac{\text{Number of patents in year } t \text{ where keyword } w \text{ appears}}{\text{Total number of patents in year } t} \tag{6}$$

This frequency metric captures the prevalence of keyword w among patents granted in year t, providing insights into the relative emphasis on AI technologies. To track changes in keyword frequency over the study period relative to the base year (2004), we define the relative frequency as:

$$Relative\ Frequency_{w,t|2004} = \frac{Frequency_{w,t}}{Frequency_{w,2004}} \tag{7}$$

Here, $Relative\ Frequency_{w,t|2004}$ normalizes the frequency of keyword w in year t against its frequency in 2004, allowing for a clear comparison of the increasing prominence of AI-related innovations over time.

Figure 5 presents the results, indicating an evident shift toward AI technologies, particularly in the later years of the study period. Analyzing keyword trends within patent documents reveals a substantial increase in the prevalence of AI-related terms, such as "deep learning", "reinforcement learning", "natural language processing", and "machine learning", which become predominant by 2019. For example, the frequency of "natural language processing" in 2015 is eight times higher than in 2004, rising to over 20 times by 2019. Likewise, "machine learning" exhibits an even more substantial increase, exceeding 30 times its 2004 frequency, underscoring the rapid expansion of AI-related innovation. Notably, the acceleration of AI-focused patenting activity becomes more pronounced after 2010, as reflected in the significant growth of keyword occurrences between 2010 and 2015.

These findings underscore an observable and accelerating shift toward AI innovation, reflecting broader transformations in the technological and intellectual property landscapes.

3.3 AI Distance and Closeness

After identifying AI patents, this section introduces two measures to quantify the proximity between any given patent and AI patents. By analyzing the dynamic evolution of this proximity over time, we can assess whether firms with increasing robot exposure progressively redirect their innovation activities toward AI-related fields. This approach provides a systematic framework to examine the technological transition induced by automation and its implications for firms' innovation strategies.

AI Distance. Comparing the distance between two patents within a patent-to-patent citation network provides an effective means of measuring their proximity. Building on this insight, we first aggregate AI patents into a singular "representative" AI patent and construct an annual patent-to-patent citation network. Within this network, two patents, p and p', are considered directly connected if a citation link exists between them, regardless of direction. The distance between two patents is defined as the number of citation links required to establish a connection, where a direct citation corresponds to a distance of one. This approach enables us to systematically capture the interconnections between AI and other technologies within the patent landscape. A patent's "AI distance" is determined by the shortest path linking it to the "representative" AI patent within the citation network. By definition, AI patents are assigned an AI distance of zero, signifying their direct association with AI technology.

AI Closeness. To further assess the relationship between individual patents and AI technology, we employ a citation-based approach to measure their "closeness" to a collective representation of AI innovation. This method begins by consolidating all AI-related patents identified via keyword analysis into a singular "representative" AI entity. For each year t, we compile the citation patterns of this representative AI patent across various USPTO technology classes, represented by the vector:

$$Cites_{AI,t} = (Cites_{AI,1,t}, \ldots, Cites_{AI,N,t}),$$

where $Cites_{AI,i,t}$ indicates the count of citations within technology class i by the AI representative at year t. To normalize these counts into a distribution, we define:

$$C_{AI,t} = (C_{AI,1,t}, \ldots, C_{AI,N,t}),\ where\ C_{AI,i,t} = \frac{Cites_{AI,i,t}}{\sum_i Cites_{AI,i,t}},$$

resulting in a citation distribution vector $C_{AI,t}$ for the AI representative, highlighting its engagement across the technological spectrum. For any given patent p filed in year t, we similarly construct a citation distribution vector $C_{p,t}$. The closeness of patent p to AI innovations—termed "AI Closeness" —is then quantified by calculating the cosine similarity between $C_{AI,t}$ and $C_{p,t}$:

$$AI\ Closeness_{p,t} = \frac{C\prime_{AI,t} C_{p,t}}{||C_{AI,t}\ ||_2 ||C_{p,t}\ ||_2}, \tag{8}$$

where $||x||_2$ denotes the L_2 norm of vector x, facilitating the normalization of the vectors for a fair comparison. Additionally, AI patents are assigned an AI Closeness score of one[10].

Overall, these two measures of AI distance and AI closeness serve as refined indicators of a patent's alignment with the forefront of AI technology, as reflected in its citation patterns. By leveraging these metrics, we can examine how firms strategically position their patenting activities toward AI, providing valuable insights into the broader trend of AI integration within the innovation landscape.

Specifically, we examine whether a firm's robot exposure influences the orientation of its patents toward AI by analyzing AI distance and AI closeness using the following specification:

$$y_{i,p,t+j} = \alpha_i + \alpha_t + \beta \times Robot\ exposure_{it} + X'_{it}\gamma + \varepsilon_{i,p,t}, \tag{9}$$

[10]One main concern regarding the AI Closeness measure is whether patents classified within the AI group exhibit high AI Closeness, considering the substantial heterogeneity among AI patents themselves. To evaluate the measure's validity, we randomly select 10% of AI patents and label them as non-AI while treating the remaining 90% as representative AI patents. We then calculate the AI Closeness of all other non-AI patents, which includes both the sampled pseudo non-AI patents and the actual non-AI patents, in relation to the representative AI patents. The findings reveal that the average AI Closeness for our randomly sampled AI patents is significantly higher than that for the true non-AI patents, with a mean difference of 0.118. Furthermore, the histogram distribution of AI Closeness for the randomly chosen AI patents typically exhibits a more right-skewed pattern compared to that of non-AI patents. This distribution suggests a more pronounced alignment between our AI Closeness measure and genuine AI innovation (Figure A2).

where $y_{i,p,t+j}$ represents either the logarithm of AI Distance plus one or the AI Closeness for patent p of firm i in year $t+j$. Robot Exposure$_{it}$, the vector $\mathbf{X}_{it}$, and the firm α_i and year α_t fixed effects are consistent with those used in Equation (5).

Table 5 presents the second-stage results of the IV estimation. Columns (1)–(3) report the IV estimates for the effect of current, one-year lagged,d and two-year lagged robot exposure on AI distance, while Columns (4)–(6) present the corresponding results for AI closeness. The findings suggest that firms with greater robot exposure tend to innovate in closer proximity to existing AI patents. Specifically, a one-standard-deviation increase in one-year lagged robot exposure is associated with a reduction in AI distance of approximately 0.581, which is statistically significant at the 1% level. Furthermore, firms with high robot exposure are more likely to produce patents with citation patterns similar to those of AI patents. On average, a one-standard-deviation increase in robot exposure leads to a 0.084 increase in AI closeness in the following year.

Besides the closeness and distance measures, as a complement, we further investigate whether the absolute number of AI patents increases. Specifically, we model AI patent grants as a time-varying Poisson process and estimate the underlying arrival rate as:

$$\mathrm{E}\left(\mathrm{y_{i,t+j}} \mid \mathbf{X_{it}}\right) = \exp\left(\alpha_i + \alpha_t + \beta \times \text{Robot Exposure}_{it} + \mathbf{X}'_{it}\gamma\right), \tag{10}$$

where $\mathrm{E}\left(\mathrm{y_{i,t+j}} \mid \mathbf{X_{it}}\right)$ denotes the expected number of AI patents filed by firm i in year $t+j$. The term Robot Exposure$_{it}$, the covariate vector $\mathbf{X}_{it}$, as well as firm and year fixed effects α_i and α_t respectively, are consistent with those used in Equation (5). Columns (7)–(9) of Table 5 present the results, which indicate that firms with high robot exposure are more likely to produce a greater number of AI-related patents. On average, a one standard deviation increase in robot exposure leads to a 0.244 increase in the number of AI patents granted in the subsequent year.[11]

[11] This result provides additional evidence supporting the hypothesis that robot exposure positively influences AI innovation within firms.

3.4 An Alternative Measure of AI Patents

To address the growing interest of policymakers and researchers, the USPTO has launched a project to compile an AI patent dataset (Giczy et al. [2022]). While a notable limitation of this dataset is its relatively low classification accuracy, it categorizes patents into AI subclasses, enabling a more granular analysis. Therefore, as a robustness check, we leverage this dataset to analyze which specific AI subclasses firms increasingly focus on in response to heightened robot exposure.

Explicitly, the USPTO first constructs a seed dataset through manual review, distinguishing AI-related patents from non-AI patents. It then applies a machine learning model, specifically the long short-term memory (LSTM) method, to classify all U.S. patents. The final dataset categorizes AI patents into eight subclasses: 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 hardware (HW). To mitigate concerns regarding classification accuracy, we aggregate patents from each AI subclass into a single "representative" AI patent for each year. We then evaluate AI Closeness and AI Distance using the methodology outlined in Section 3.3.

Table 6 presents the results, showing that firms with higher robot exposure increasingly direct their innovation efforts toward AI across all subclasses, as evidenced by higher AI closeness. In Panel A, firms with greater exposure to robotics exhibit significantly higher AI Closeness, with all coefficients being both economically and statistically meaningful. Panel B examines AI Distance, where the estimated coefficients are predominantly negative. Notably, three subclasses—evolutionary computation, natural language processing, and hardware—display significantly negative coefficients. This finding suggests that firms with higher robot exposure tend to be closer to the technological frontier within each AI subclass in the citation network.

3.5 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 et al. [2022], which demonstrate that government procurements in China significantly stimulate AI innova-

tion by granting access to proprietary data, we investigate whether firms with substantial robot exposure similarly gain access to richer data resources. This improved data availability may accelerate AI advancements, 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 reinforcing cycle of technological progress.

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.

To examine whether increased robot exposure leads to greater data availability, we must quantify the extent of data accessible to firms, which presents a significant challenge. Following an approach similar to that used for measuring 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 A7, are derived from a broad body of literature [Chen et al., 2012, Farboodi et al., 2019, McAfee et al., 2012, Brynjolfsson and McElheran, 2016, Brynjolfsson and Mitchell, 2017, Brynjolfsson et al., 2010, Agrawal et al., 2019, Begenau et al., 2018, Tambe, 2014, Zhu, 2019, Goldstein et al., 2019, Chen et al., 2019, Fuster et al., 2019, Bajari et al., 2019, Wu et al.,

2020]. Our methodology involves computing the frequency of these keywords in firms' 10-K annual reports and normalizing them by industry-specific frequencies to ensure comparability across sectors. Formally, we define the data generation measure as follows:

$$\text{Data Generation}_{i,t} = \sum_{w} \frac{\text{Data Word Number}_{w,i,t}}{\text{Data Word Number}_{w,t}}, \tag{11}$$

where $\text{Data Word Number}_{w,i,t}$ represents the count of data-related keywords w appearing in firm i's 10-K filing during year t. This approach provides a standardized measure of a firm's engagement with data generation, allowing us to assess its relationship with robot exposure.

Table 7 presents evidence that firms with greater robot exposure experience a significant increase in data generation. Columns (1)–(3) report the average effects of rising robot exposure on data generation in the current year and the subsequent two years. The results indicate that a one-standard-deviation increase in one-year lagged robot exposure is associated with an increase in data generation of approximately 1.428, which is statistically significant at the 5% level. Columns (4)–(5) further examine heterogeneity across firm size. Specifically, we categorize firms into two groups based on sales and net income: small firms (bottom 75%) and large firms (top 25%). The results reveal that the interaction terms between robot exposure and the large-firm indicator (for both sales and net income) are significantly positive. These findings suggest that larger firms derive the greatest benefits, exhibiting the most substantial increase in data generation as robot exposure rises. This result is consistent with the non-rivalrous nature of data, wherein larger firms are better positioned to leverage data-driven advancements.

4 Aggregate Innovation

In the preceding section, we established that increasing automation prompts firms to reorient their innovation focus toward AI-related fields. Nonetheless, 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.

We approach the aggregate implications from two perspectives: patent innovation (output) and R&D (input). On the output side, we find evidence of an innovation "J-curve," characterized by an initial decline followed by a subsequent increase in the citation-weighted number of patents. This pattern suggests that while firms may experience a temporary drop in patent output as they adjust to the integration of automation, there is a recovery phase where innovation flourishes. On the input side, we observe a significant rise in R&D expenditure in response to increasing robot exposure, indicating that firms are willing to invest more in innovation when they adopt new automation technologies.

The contrasting pattern observed in the short term—a rise in R&D spending concurrent with a decline in patent innovation—merits further exploration. We provide further evidence to show that several forces may underlie this phenomenon: (i) Firms might be shifting toward more foundational research, which is characterized by higher levels of originality and generality; (ii) AI-related innovations could entail greater incremental costs compared to non-AI patents, due to factors such as larger research teams, a higher labor input, and greater inventor originality. These investments could ultimately lead to higher-value innovations; (iii) Significant investments in human capital accumulation are critical, as AI-driven innovation requires specialized expertise and a skilled workforce capable of working on the frontiers of technology. Bloom et al. [2020] documents that innovation productivity has declined recently through case studies, and we provide one potential micro channel driving this declining innovation productivity – innovation shift.

These findings suggest a nuanced relationship between automation and innovation, where initial investment and transition costs may lead to a temporary reduction in patent output, but the potential for high-value innovation output exists in the long term. The "J-curve" effect in innovation offers a compelling narrative for how firms may navigate the complexities of integrating advanced technologies such as AI and robotics into their strategic innovation efforts.

4.1 Innovation "J-Curve"

To quantify innovation outputs, we consider two key metrics: patent count and patent value.

Regarding patent count, following Cohn et al. [2022], we model the patent grant process as a time-varying Poisson process, with the underlying arrival rate specified as:

$$\mathbb{E}\left(y_{i,t+j} \mid \mathbf{X}_{it}\right) = \exp\left(\alpha_i + \alpha_t + \beta \times \text{Robot Exposure}_{it} + \mathbf{X}'_{it}\gamma\right), \tag{12}$$

where $\mathbb{E}\left(y_{i,t+j} \mid \mathbf{X}_{it}\right)$ represents the expected citation-weighted patents filed by firm i in year $t+j$, following the method proposed by Acemoglu et al. [2016]. Specifically, we calculate the citations that each patent received in the five years following its application and adjust for citation inflation by normalizing the average citations within each three-digit CPC class. The vector $\mathbf{X}_{it}$, and the firm α_i and year α_t fixed effects are consistent with those used in Equation (5).

For patent value, we estimate the following equation:

$$y_{i,t+j} = \alpha_i + \alpha_t + \beta \times \text{Robot Exposure}_{it} + \mathbf{X}'_{it}\gamma + \varepsilon_{i,t}, \tag{13}$$

where $y_{i,t+j}$ represents the citation-weighted patent value of firm i in year $t+j$, measured following the methodology proposed by Kogan et al. [2017].

The results of our analysis are presented in Table 8. Columns (1)–(3) report the estimated coefficients for robot exposure in year t on the citation-adjusted number of patents in years t, $t+1$, and $t+2$, respectively. The findings reveal a significant negative association between increased robot exposure and patent innovation in year t. Specifically, a one-standard-deviation increase in robot exposure corresponds to a reduction in patent count of approximately 0.086, which is statistically significant at the 5% level.

To account for heterogeneity in patent value, Columns (4)–(6) present the second-stage 2SLS estimates for citation-adjusted patent value. The results indicate that higher levels of robot exposure are significantly associated with declines in patent value, with estimated reductions of 0.997, 1.165, and 1.601 in years t, $t+1$, and $t+2$, respectively. Overall, these findings suggest that rising robot exposure is linked to a substantial decline in patent production, irrespective of whether adjustments for patent value are considered.

We further investigate the persistence of the decline in patent activity following increased robot exposure by examining the cumulative number and value of patents over a six-year hori-

zon. Table 9 presents the results, where patent activity is measured using both citation-weighted patent count and patent value. The findings reveal that as robot exposure rises, the cumulative number of patents experiences a significant decline during the first two years, with the magnitude of this decline gradually diminishing over time. By the fifth year, the coefficient associated with robot exposure shifts from negative to positive, leading to a notable increase in the number of patents in the sixth year. This pattern follows a "J-curve" trajectory, wherein patent output initially decreases before experiencing a subsequent rebound.

Additionally, the decline in cumulative patent value remains statistically significant only during the first four years, with the coefficient exhibiting a downward trend. After this period, the effect becomes statistically insignificant. These findings suggest that the impact of robot exposure on innovation outputs is dynamic and time-dependent, initially suppressing patent activity before facilitating a recovery in later years.

Beyond analyzing overall trends, we examine which specific patent classes experience a decline and whether any exhibit an increase. To do so, we classify patents based on their two-digit CPC codes and introduce an interaction term between these categories and robot exposure in Equation (12). Figure 6 presents the top 10 CPC industries with significantly negative interaction term coefficients, along with those industries where the interaction term is significantly positive. The results indicate that most CPC patent categories experience a decline in patent numbers following increased robot exposure. However, in the CPC categories where patent numbers are increasing, several are closely associated with AI innovation, including mechanical components, amusements, and checking devices. Given that the CPC classification does not explicitly define AI patents, the identification of AI-related patents, as outlined in the previous section, remains essential for a more precise analysis.

Take together, our findings demonstrate that increased robot exposure initially suppresses firm-level innovation, reducing both patent count and value. However, this effect follows a "J-curve" pattern, with patent output recovering after an initial decline. These results highlight the dynamic nature of automation's impact, suggesting that robot exposure ultimately fosters technological advancement, aligning with the findings of Babina et al. [2024], who document that AI investments, though slow-moving, enhance long-term innovation rather than generating immediate effects.

4.2 Rising R&D Expenditure

Next, we examine whether firms adjust their R&D expenditures in response to robot exposure. Specifically, we analyze the impact of robot exposure on R&D intensity using the following estimation model:

$$y_{i,t+j} = a_i + a_t + \beta \times Robot\ exposure_{i,t} + X'_{it}\gamma + \varepsilon_{i,t} \tag{14}$$

where $y_{i,t+j}$ represents firm i's R&D intensity, measured as R&D expenditures scaled by total assets or sales in year $t+j$. The key independent variable, Robot Exposure$_{it}$, along with the control vector $\mathbf{X}_{it}$ and firm α_i and year α_t fixed effects, are specified consistently with those in Equation (5).

Table 10 presents the second-stage IV estimation results. Columns (1)–(6) report the coefficients of robot exposure on R&D expenditure, normalized by total assets and total sales, respectively. A one standard deviation increase in robot exposure is associated with a 0.207 increase in R&D-to-total-assets and a 0.684 increase in R&D-to-total-sales in the following year. These effects are both economically and statistically significant at the 5% level. 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.

4.3 Generality and Originality

A plausible explanation for the observed innovation "J-curve" is that the shift toward 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 originality and generality. 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 patenting less. However, this shift may lead to fewer but higher-quality patents with greater technological significance.

In our context, firms' transition toward AI innovation, particularly in cutting-edge technologies such as deep learning and reinforcement learning, often necessitates breakthroughs in foundational science and technology. Advancements in AI rely on a 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. As a result, AI-driven patent innovation compels firms to revisit fundamental challenges, striving for original breakthroughs in AI theories and methodologies that surpass existing technological constraints. Consequently, firms may increase their R&D investment, but the changing composition of R&D activities—with a stronger emphasis on fundamental research—could explain the decline in patenting, as firms prioritize scientific advancements over immediate commercialization.

To empirically verify whether firms with greater automation exposure focus more on fundamental research, we adopt the approach outlined by Hall et al. [2001] to measure a patent's originality and generality:

$$Generality_p = 1 - \sum_{j=1}^{N} s_{pj}^2, \; Originality_p = 1 - \sum_{j=1}^{N} q_{pj}^2 \tag{15}$$

where s_{pj} represents the share of citations that patent p receives from technology class j within five years of its issuance, while q_{pj} denotes the share of citations patent p makes to technology class j. Empirically, we define technology classes 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 technological impact beyond its immediate 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 issue, we measure a patent's generality based on citations received within the five years following its filing year t, thereby mitigating biases stemming from long-term citation trends.

We estimate the following specification:

$$y_{i,p,t+j} = \alpha_i + \alpha_t + \alpha_{c,t} + \beta \times Robot\ Exposure_{i,t} + X'_{it}\gamma + \varepsilon_{ipt}, \tag{16}$$

where $y_{i,p,t+j}$ represents the originality or generality measure of patent p, filed in year $t+j$ by firm i. To account for potential time-varying technological factors, we control for a patent's filing year by its technology class fixed effect ($\alpha_{c,t}$), where c denotes the technology class of patent p. This specification helps mitigate biases arising from technological heterogeneity across time. The key independent variable, Robot Exposure$_{it}$, along with the control vector $\mathbf{X}_{it}$ and firm (α_i) and year (α_t) fixed effects, are defined consistently with those in Equation (5).

Table 12 presents the results. Columns (1)–(3) and (4)–(6) show that higher robot exposure is associated with increases in both generality and originality of innovation in the current and subsequent years. Specifically, a one standard deviation increase in robot exposure leads to an approximate 0.022 increase in patent generality in year $t+1$, which is both economically significant (given a mean generality of 0.128) and statistically significant at the 1% level. Likewise, a one standard deviation increase in robot exposure results in an approximate 0.018 increase in patent originality over the same period.

These findings suggest that during the transition towards AI-driven innovation, firms place greater emphasis on foundational research characterized by higher originality and generality. 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 by studies such as Arora et al. [2018], a key transformation since the 1980s has been the reallocation of resources and attention by leading firms away from exploratory scientific research toward more commercially oriented projects. However, our research identifies an opposing trend, wherein the shift toward foundational AI innovation can, in turn, stimulate more exploratory scientific research.

4.4 The Relative Cost of AI-Driven Innovation

Beyond foundational research, the rising relative costs associated with AI-driven innovation may also contribute to the observed increase in R&D expenditure and the initial decline in patent output. This section examines the incremental costs of AI patents relative to non-AI patents across multiple dimensions. Building on the theoretical framework proposed by Jones [2009], we construct several measures to quantify the costs associated with innovation. Specifically, we assess the complexity and resource intensity of AI innovation by considering factors such as team size, labor input, inventor originality, and patent value. These dimensions help capture the broader investment requirements and structural differences that distinguish AI-related innovation from traditional technological advancements.

The first two measures focus on inventors. First, 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 Jones [2009]. 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. Second, 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. This adjustment provides a more accurate representation of the effective labor effort allocated to each innovation project, mitigating potential overestimation caused by overlapping contributions.

Third, prior research Hall et al. [2001], Jones [2009] suggests 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 innova-

tion process. We quantify an inventor's originality using the following measure:

$$Inventor\ Originality_{i,t} = 1 - \sum_{j=1}^{N} s_{ji,t}^2 \tag{17}$$

where $s_{ji,t}$ represents the fraction of patents by inventor i that fall into technology class j over the past five years. A higher originality score indicates that an inventor has contributed to innovations across a more diverse set of technological domains. For each patent, we compute its inventor originality as the average originality score of all inventors associated with the patent. This measure provides insight into the breadth of technological exploration undertaken by the inventors and serves as a proxy for the cognitive and interdisciplinary demands of AI-driven innovation.

Finally, 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, substantial 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 structure.

Figure 7 presents a comparative analysis of AI and non-AI patents across key cost-related dimensions. Subfigure (a) displays the average team size of AI and non-AI patents over time, showing that AI patents consistently involve larger inventor teams than non-AI patents. The gap between the two groups widened significantly after 2015, aligning with the increasing shift toward AI-driven innovation. Subfigure (b) illustrates the time trend of average labor input, revealing that AI patents consistently require greater labor input than non-AI patents. This pattern suggests that AI-related innovations demand higher resource investments, reflecting their greater complexity and development costs. Subfigure (c) depicts the average inventor originality, indicating that AI patents generally involve inventors with broader technological expertise. While AI patents have historically exhibited higher originality than non-AI patents, the gap has narrowed in recent years, possibly due to the expanding integration of AI across various technological domains. Subfigure (d) presents the average patent value, demonstrating that AI patents hold consistently higher economic significance than non-AI patents. The divergence became

particularly pronounced after 2011, suggesting that AI innovations yield patents of greater technological impact and commercial value. Together, these findings underscore the distinct cost structures associated with AI-related innovations and highlight their increasing prominence in the technological landscape.

We further validate these findings using an OLS regression approach. Table A8 presents the regression results corresponding to the trends depicted in Figure 7. Even after controlling for firm-year fixed effects, the results confirm that AI patents consistently incur higher costs relative to non-AI patents across multiple dimensions. Specifically, AI patents are associated with larger inventor teams, greater labor input, higher inventor originality, and increased patent value. These findings reinforce the notion that AI-related innovations demand more substantial resource investments, reflecting their inherent complexity and the need for interdisciplinary collaboration.

Overall, the regression analysis provides robust empirical support for the observed cost differentials between AI and non-AI patents, further highlighting the distinctive nature of AI-driven innovation.

4.5 Standing on Jones' Shoulders

The third possible explanation for the innovation "J-curve" is the substantial 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 Yang and Zhu [2020]. 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 greater prior AI experience are likely to have developed specialized expertise and accumulated knowledge, which may facilitate a smoother and more efficient transition toward AI-driven innovation. By 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.

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:

$$\begin{aligned} y_{i,p,t} = \alpha_i + \alpha_t + \beta_0 \times \textit{Robot exposure}_{it} + \beta_1 \times \textit{AI experience}_{i,03} \times \textit{Robot exposure}_{it} \\ + X'_{it}\gamma + \varepsilon_{i,p,t}, \end{aligned} \tag{18}$$

where $y_{i,p,t}$ represents the AI distance (or closeness) of firm i's patent p in year t. $\textit{AI experience}_{i,03}$, is a dummy equal to 1 if firm i's AI experience prior to 2004 exceeds the industry median and 0 otherwise. The variable Robot Exposure$_{it}$, along with the control vector $\mathbf{X}_{it}$ and firm (α_i) and year (α_t) fixed effects, are defined consistently with those in Equation (5). If prior research experience facilitates the transition toward AI-related innovations, we expect firms with greater AI research experience to exhibit a more pronounced shift toward AI fields in response to rising automation compared to their less experienced counterparts.

Table 11 presents the heterogeneous effects of rising robot exposure on the shift toward AI innovation, measured by AI distance and AI closeness. Columns (1) and (2) report the estimation results, revealing two key patterns. First, increased robot exposure is associated with a significant reduction in AI distance and an increase in AI closeness for both groups of firms. Second, the magnitude of this shift is nearly twice as large for firms with substantial prior AI experience, suggesting that firms with a stronger AI foundation are better positioned to integrate automation-driven advancements into their innovation strategies. These findings highlight the shift costs in shaping firms' adaptive responses to automation, reinforcing the importance of accumulated knowledge in facilitating technological transitions.

4.6 The Persistence of Innovation Shift

In addition to the long-term effects of robot exposure on innovation productivity discussed earlier, we also examine its prolonged impact on the shift in innovation focus. Table A9 investigates the persistence of the effects of increasing robot exposure over the subsequent ten years. The results indicate that the transition toward AI-driven innovation exhibits significant long-term effects, with AI Closeness remaining significantly positive even after a six-period lag.

This finding suggests that firms continue to deepen their engagement with AI-related technologies well beyond the initial exposure to automation. In contrast, the effects of robot exposure on other innovation-related variables appear to be more transient, persisting for a maximum of four lagged periods. This pattern implies that while general innovation responses to automation may be temporary, the strategic shift toward AI innovation represents a more enduring transformation.

5 Robustness Checks

In this section, we conduct a series of robustness checks to validate the core findings on innovation shift and innovation productivity. To ensure the reliability of our results, we employ multiple approaches, including assessing the validity of IV estimates, constructing alternative measures of robot exposure, examining the effects of the Supreme Court's Alice decision, exploring alternative sample constructions, and analyzing results at different levels of aggregation. These robustness tests help confirm the consistency and robustness of our empirical findings.

5.1 Assessing the Validity of IV Estimates

5.1.1 Evaluating Common Shock Effects and Pre-trend Tests

The validity of using robot usage per worker in 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 enhance the robustness of our findings, we implement two robustness checks to address these concerns.

To account for potential common shock effects, we consider industry-level price changes as a possible channel, as fluctuations in industry prices can create economic pressures or incentives that influence firms' investment decisions, including automation technologies such as robots, as

well as their innovation strategies. To address this concern, we obtain the Producer Price Index (PPI) at the industry level from the U.S. Bureau of Labor Statistics, which provides annual variations. We incorporate this variable into the regression model to assess whether its inclusion affects the estimated results. Additionally, we control for industry-year fixed effects in the regression, where industry is defined at the SIC 2-digit level, to better capture industry-specific factors that vary over time.

The regression results, presented in Tables A10 and A11, indicate that after incorporating these additional controls, the estimated results remain largely unchanged. This finding suggests that the baseline results are robust to potential confounding effects arising from industry-level price fluctuations and other industry-specific shocks.

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 to test this hypothesis.

First, as a placebo test, we examine whether robot exposure during the 2004–2019 period influences firms' innovation shifts and innovation productivity between 1988 and 2003. The regression results, presented in Table A12, indicate that the estimated coefficients for innovation shifts and inputs are not statistically significant. Moreover, the direction of innovation in the earlier period reverses, suggesting that before 2004, firms with high levels of robot adoption do not exhibit a positive trend toward AI-related innovation.

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. Table A13 shows that incorporating this control does not affect the parameter estimates from the baseline results.

These findings confirm that no pre-existing trends bias our baseline regression outcomes, reinforcing the validity of our empirical strategy.

5.1.2 Alternative IV Construction

To assess the robustness of our baseline results, we employ two alternative methods to construct the instrumental variable.

The first method expands the robot usage data to include additional EU countries. In our baseline analysis, we construct the IV using robot adoption data from five European countries—Finland, Sweden, Denmark, Italy, and France. However, this selection may not fully capture broader automation trends across Europe. To assess the robustness of our findings, we modify the industry-level instrument in two ways, following the procedure in Acemoglu and Restrepo [2020]: (1) by incorporating Germany into the five-country average, and (2) by expanding the average to include all nine advanced automation countries in Europe: Finland, Sweden, Denmark, Italy, France, Germany, Norway, Spain, and the United Kingdom.[12] These nine countries are selected due to their comprehensive industry statistics available in the IFR database and their combined representation of 41% of the global industrial robot market. As noted by Acemoglu and Restrepo [2020], robot adoption rates in the five baseline countries closely track those in the United States, whereas Germany's adoption rate is significantly higher. Meanwhile, Norway, Spain, and the United Kingdom exhibit adoption rates comparable to those of the United States.

In the second approach, we account for potential time-related trends or unobserved time-specific factors that may influence the IV. To do so, we regress the IV on time fixed effects, thereby isolating any systematic variation attributable to temporal dynamics. We then extract the residuals from this regression and use them to re-estimate the main results. This adjustment ensures that the instrument is not confounded by broader macroeconomic trends or other time-dependent factors that could bias the estimates.

The results, presented in Tables A14, A15, and A16, confirm the robustness of our findings. Across all alternative constructions, the conclusions drawn from the baseline analysis remain unchanged. These results reinforce the validity of our IV strategy, demonstrating that the observed effects are not driven by the selection of European countries or spurious time-related fluctuations.

[12] Since data for Norway are unavailable in the EU KLEMS, we impute the Norwegian distribution using employment data from Denmark, Finland, and Sweden.

5.1.3 Construction of the Shift-Share IV

To mitigate the potential issue that both robot exposure in 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 the shift-share instrumental variable (SSIV) approach in this section. Specifically, we assume that robot usage per worker across EU countries serves as a quasi-experimental shock, as proposed by Borusyak et al. [2022]. The rationale behind this approach builds on the observation by Acemoglu and Restrepo [2022] that differences in industry-specific aging patterns in earlier years led to varying degrees of 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. This methodology ensures that the instrument captures exogenous variation in robot exposure while mitigating potential endogeneity concerns.

Specifically, we modify the original IV by replacing its time-varying component—representing the annual proportion of robot usage to employment across EU countries—with a fixed average proportion calculated over the period 2000 to 2004, the five years preceding the start of the sample. The construction of this IV is as follows:

$$Robot\ Exposure\ IV_{i,t} = \sum_{k=1}^{m} \frac{Sales_{i,j,t}}{Sales_{i,t}} \times \frac{Robots^{EU}_{j,\text{00-04}}}{Workers^{EU}_{j,\text{00-04}}}, \tag{19}$$

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

Before applying the SSIV approach, we conduct a series of validity tests to ensure that the exogenous shocks satisfy the requirements of quasi-randomness and large-sample consistency, thereby enabling reliable SSIV estimates. To assess sample adequacy, we calculate the effective sample size across SIC15 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. According to Borusyak et al. [2022], even an effective sample size as small as

20 can yield reasonably accurate statistical inference in industry-level regressions. Additionally, we evaluate the within-group correlation of shocks under the SIC15 classification and confirm that the condition of sufficiently uncorrelated shocks is met, further supporting the validity of our SSIV estimation.

We further 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.[13] Following Acemoglu and Restrepo [2020], we include a set of industry-level control variables, such as the log of real wages (Ln(Wage)), wage per worker (Wage/Worker), the log of employment (Ln(Worker)), the share of high-tech equipment in total investment (High-tech/Invest), the ratio of input to value-added (Input/Value-added), and the ratio of workers to value-added (Worker/Value-added). At the firm level, we incorporate the baseline control variables from the study along with a dummy variable indicating whether the firm employs an ES system. Additionally, we perform a firm-level pre-trend analysis by regressing pre-trend variables from 1988 to 2003 on the shift-share instrument. If the shocks are as-good-as-randomly assigned to industries within periods, they should not predict the predetermined variables. The results in Table A17 confirm that there are no statistically significant relationships between these variables and the shift-share instrument within periods, except for firm age. To address this issue, we conduct a sensitivity test to evaluate the robustness and stability of the results, ensuring that our findings remain consistent despite potential confounding effects.

Table A18 presents the SSIV estimates from regressions of innovation shifts and inputs in the U.S. on robot exposure, instrumented by robot usage across EU countries. Following Borusyak et al. [2022], we estimate the coefficients using equivalent industry-level regressions to obtain valid exposure-robust standard errors. Consistent with the methodology outlined above, we cluster standard errors at the SIC15 level to account for potential within-industry correlation. The results in Table A18 closely align with the main findings presented earlier, confirming that increased robot exposure leads to a shift in innovation toward AI and higher R&D investment.

[13] Following Borusyak et al. [2022], we estimate the coefficients using equivalent industry-level regressions to obtain valid exposure-robust standard errors.

5.2 Alternative Robot Exposure Measures

We explore alternative approaches to measuring robot exposure. In the first method, we refine and reweight keywords using a two-stage regression process, following Cong et al. [2019]. The first stage employs a LASSO regression for variable selection:

$$y_{jt} = \sum_{w \in W} \beta_w x_{w,j,t} + \varepsilon_{jt} \tag{20}$$

where y_{jt} represents robot exposure in industry j for year t, defined as the number of industrial robots installed per thousand workers, and W denotes the set of keywords. In the second stage, we identify keywords with non-zero coefficients from the LASSO regression and use them in an OLS regression to refine the coefficient estimates. The refined subset of keywords is denoted as W^*, with their corresponding OLS coefficients represented by $\beta_w, w \in W^*$. Using this refined keyword set, we compute each firm's robot exposure for year t as follows:

$$\text{Robot Exposure}_{it} = \sum_{w \in W^*} \beta_w x_{i,w,t}, \tag{21}$$

where $x_{i,w,t}$ represents the scaled occurrence of keyword w in firm i's 10-K filing for year t.

In the second approach, we measure robot exposure using the ratio of robots to employment. Robot usage data is obtained from the IFR, while employment data is sourced from EU KLEMS, following the methodology outlined in Acemoglu and Restrepo [2020]. To construct firm-level robot exposure, we integrate the industry-level robot exposure dataset with the Compustat-Segment dataset, which provides detailed information on firms' business activities across industries. Firm-level robot exposure is defined as:

$$Robot\ Exposure_{i,t} = \sum_{j=1}^{m} \frac{Sales_{i,j,t}}{Sales_{i,t}} * \frac{Robots_{j,t}}{Workers_{j,t}} \tag{22}$$

where $Sale_{i,j,t}$ represents the annual sales of firm i in segment (industry) j at year t, and $Sale_{i,t}$ denotes the total sales of firm i in year t. The term $\frac{Robots_{j,t}}{Workers_{j,t}}$ captures industry-level robot exposure for industry j at year t.

Third, a potential concern is that our empirical findings may be driven by time trends or

co-integration rather than the causal relationship between robot exposure and innovation outcomes. To address this issue, we first regress robot exposure on time fixed effects and use the residuals from this regression as the key explanatory variable. This approach effectively filters out any time-specific patterns, ensuring that the variation in robot exposure is not confounded by overarching temporal trends.

Finally, since robots, as capital assets, depreciate over time, relying on data from a single year may not accurately capture their actual value and utilization during that period. To address this issue, rather than using a one-year mention in 10-K filings, we adopt a rolling five-year weighted average to measure robot exposure. This approach accounts for depreciation by applying an annual discount rate of 0.8, ensuring that past investments in robotics continue to influence the firm's exposure measure while gradually declining in weight over time. By incorporating a rolling average, this method smooths fluctuations in robot valuation and usage over their lifecycle, mitigating potential misinterpretations caused by year-to-year volatility.

Tables A19, A20, A21, and A22 present the results from these alternative measurement approaches. The magnitude and statistical significance of the core variable coefficients remain consistent with the baseline findings, reinforcing the robustness of our results.

5.3 Effects of the Supreme Court's Alice Decision

This subsection examines the influence of the Supreme Court's 2014 *Alice Corp. v.s. CLS Bank International* decision on AI patentability. The *Alice* ruling significantly impacted patent law, particularly regarding software, business methods, and abstract concepts. The Court held that inventions based on "abstract ideas" are not patentable unless they include "sufficient inventiveness" or "additional inventive concepts." As a result, AI innovations involving abstract algorithms or processes may face increased difficulty in securing patent protection post-*Alice*.

First, we assess whether the *Alice* decision has made it harder to obtain AI patents. Using USPTO pre-grant data, we identify AI patents through abstract analysis and merge these with our dataset of granted patents, covering applications from 2004 to 2019. We measure the impact using two variables: $Dummy(Grant)$, a binary indicator of whether a patent is granted, and $Ln(Grant\ Lag)$, the logarithm of the days from application to approval (plus one).

To test differences in approval rates, 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.

Table A23 summarizes the results. Columns (1) and (2) show that the interaction term's coefficient is negative but not statistically significant, indicating no substantial impact of *Alice* on AI patent approval rates. Column (3) examines 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 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.

Next, we examine the *Alice* decision's impact on innovation shifts, inputs, and outputs by including an interaction term, $Robot\ Exposure \times Dummy(After\ 2014)$, in our regressions. Table A24 presents these results. In Panel A, the interaction term's coefficient is negative in the regression on technology similarity. Panel B shows a negative coefficient for AI Distance but a positive one for AI Closeness, suggesting firms exposed to robotics continue innovating in AI technologies post-*Alice*. Panels C and D indicate that this innovation shift is accompanied by increased R&D spending and decreased patent output. Overall, these results confirm that firms with greater robot exposure maintain their focus on AI innovation despite the *Alice* decision.

Finally, we conduct robustness checks whether the *Alice* case affect the AI shift, in the patent-firm-year regressions. Table A25 shows that the pass of the Alice does not significantly alter results for AI-directed innovation or innovation quality, confirming the robustness of our findings. Overall, the AI shift should not be driven by the hurdle imposed by the *Alice* case.

5.4 Alternative Approaches to Sample Construction

Since some new technologies, such as big data or cloud computing, are closely related to artificial intelligence, they might also influence our main regression results. To control for the effects of other technological shocks, we follow the methodology of Ewens et al. (2018). Firms operating in the internet industry are likely to be impacted by cloud computing and other in-

ternet technologies. Therefore, we exclude firms in the internet industry from our analysis. Additionally, following the approach of Jia et al. (2020) and based on Reisinger's (2014) statement, we focus on Enterprise Systems (hereinafter referred to as ES), which represent the most comprehensive category of information technology investments. ES is currently the most comprehensive enterprise information system globally, and its advanced technology and powerful functionality can directly influence the firm's information environment (Dorantes, 2013). As such, we use the adoption of ES as a measure of firms' information technology. To determine the ES sample, we first extracted records of ES usage from the Lexis-Nexis database using ES-related search keywords. Based on this, we excluded firms that had implemented ES systems. The results, presented in Table A26 and Table A27, demonstrate that our primary conclusions remain robust after these adjustments.

In this section, we consider two alternative sample constructions: (1) In our sample, many U.S. firms have overseas production units, which could significantly affect our main conclusions. Thus, in this subsection, we exclude firms whose overseas sales account for more than 50% of total sales and re-run the tests. (2) Given that extreme automation levels in certain firms or industries could affect the generalizability of our baseline results, we exclude the top 1% of firms in terms of automation index for each industry-year. The results of the robustness checks based on these sample adjustments are shown in Table A28 and Table A29, which remain consistent with the baseline regression results.

5.5 Citation-weighted Measure

In the previous analysis, we employed patent-level variables to more accurately examine firms' innovation transitions towards AI and improvements in patent quality. In this subsection, in order to maintain consistency with other results, we use a citation-weighted measure to calculate firm-level variables. Table A30 presents these results, which are largely consistent with the findings in the original study.

6 An Equilibrium Model of Innovation Shift

This section develops a dynamic equilibrium model to rationalize the observed innovation shift in response to rising robot exposure. Appendix 7 contains the proofs for all propositions.

6.1 Production

Consider a continuous-time economy with $t \in [0,\infty)$, consisting of a continuum of firms indexed by $f \in [0,1]$. Following the framework of 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 in the interval $s \in [\theta_f,1]$ 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 productivity. Innovations related to automated tasks are integrated with machines to produce the final product:

$$\frac{1}{1-\beta} q_M(f)^{\beta} k_M^{1-\beta}(s \mid f), \tag{23}$$

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

$$\frac{1}{1-\beta} q_L^{\beta}(f) k_L^{1-\beta}(s \mid f), \tag{24}$$

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 $\frac{1}{\chi_L \bar{q}_L}$, where

$$\bar{q}_L = \int_0^1 (1-\theta_f) q_L(f) df. \tag{25}$$

We normalize the production of $k_L(s|f)$ by $\bar{q}_L$ to ensure that the stationary equilibrium remains linear, following Akcigit and Kerr [2018].

We normalize the price of the final product to one for each time $t \in [0,\infty)$. At time t, firm f chooses $k_M(s|f)$ and $k_L(s|f)$ to maximize instantaneous profit conditional on the latest technology,

$$\begin{aligned}\Pi_f(q_M, q_L) = \max_{k_M(s|f), k_L(s|f)} \frac{1}{1-\beta} \int_0^{\theta_f} q_M^{\beta}(f) k_M^{1-\beta}(s \mid f)\, ds - \frac{1}{\chi_M} \int_0^{\theta_f} k_M(s \mid f) ds \\ + \frac{1}{1-\beta} \int_{\theta_f}^1 q_L^{\beta}(f) k_L^{1-\beta}(s \mid f) ds - \int_{\theta_f}^1 \frac{w k_L(s \mid f)}{\bar{q}_L \chi_L} ds,\end{aligned} \tag{26}$$

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 $q_L(f)$:

$$\Pi_f = \alpha_M \theta_f q_M(f) + \alpha_L (1-\theta_f) q_L(f), \tag{27}$$

where

$$\alpha_M = \frac{\beta}{1-\beta} \chi_M^{\frac{1-\beta}{\beta}}, \quad \text{and} \quad \alpha_L = \frac{\beta}{1-\beta} \chi_L^{1-\beta}. \tag{28}$$

Intuition. The marginal instantaneous return of innovation in automated tasks is given by $\frac{\partial \Pi_f}{\partial q_M(f)} = \alpha_M \theta_f := \gamma$, 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.

Put together, γ captures the total marginal return of innovation in automated tasks, which is the product of the extensive and intensive margins. In our theoretical framework, we assume that α_M is uniform across firms, while firms differ in their automation degree θ_f. Empirically, various factors can lead to changes in α_M. For instance, a rising cost in materials that complement the use of robots in production lines can be interpreted as a decline in α_M.

6.2 Innovation

Innovations in Automated Tasks. We follow the innovation framework proposed by Akcigit and Kerr [2018], modeling innovation as an arrival process. Specifically, for the technology used in automated tasks, firm f can invest in R&D to enhance the probability of technological advancements. The R&D expenditure associated with the arrival rate z_M is given by

$$C_M(z_M, q_M) = C_M(z_M) q_M, \tag{29}$$

where $C_M(0) = 0, C'_M(0) = 0, C''_M(0) > 0$, and there exists an upper bound $\bar{z}$ such that $C_M(\bar{z}) = \infty$.

To emphasize the time-varying nature of technology, we relabel $q_M(f)$ and $q_L(f)$ as $q_M(t|f)$ and $q_L(t|f)$, respectively. The notation $q_k(t|f)$ and $q_k(f)$, for $k \in \{M, L\}$, can be used interchangeably without loss of clarity.

Given the arrival rate z_M, firm f patents one innovation with probability $z_M \Delta t$ and experiences no innovation with probability $1 - z_M \Delta t$ over the interval $[t, t + \Delta t]$. The technology level of firm f at time $t + \Delta t$ is

$$q_M(t + \Delta t \mid f) = \begin{cases} q_M(t \mid f)(1 + \lambda_M) & \text{w.p. } z_M \Delta t \\ q_M(t \mid f) & \text{w.p. } 1 - z_M \Delta t \end{cases} \tag{30}$$

where λ_M is the step size of the technology process from one innovation. If there is no innovation between t and $t + \Delta$, the technology at $t + \Delta$ remains the same as at t. Otherwise, a single innovation occurs, raising the technology level by λ_M. The probability of two or more innovations within Δt is of order $o(\Delta t)$, which is negligible in the continuous-time framework.

Innovations in Non-Automated Tasks. Firm f also engages in R&D to enhance the technology used in non-automated tasks, following a similar arrival process as in automated tasks. The R&D expenditure associated with the arrival rate z_L is given by

$$C_L(z_L, q_L) = C_L(z_L) q_L, \tag{31}$$

where $C_L(0) = 0$, $C_L'(0) = 0$, $C_L''(0) > 0$, and there exists an upper bound $\bar{z}$ such that $C_L(\bar{z}) = \infty$.

Moreover, the technological progress in non-automated tasks between t and $t + \Delta t$ follows:

$$q_L(t + \Delta t \mid f) = \begin{cases} q_L(t \mid f)(1 + \lambda_L) & \text{w.p.} \quad z_L \Delta t \\ q_L(t \mid f) & \text{w.p.} \quad 1 - z_L \Delta t \end{cases} \tag{32}$$

where λ_L represents the step size of the technological improvement from one innovation. If no innovation occurs between t and $t + \Delta t$, the technology remains unchanged. Otherwise, a single innovation raises the technology level by λ_L. Similarly, the probability of multiple innovations within Δt is of order $o(\Delta t)$, making it negligible in the continuous-time framework.

6.3 R&D Decision

For a firm f, denote its value as V, which depends on q_M and q_L. Firm f solves a dynamic optimization problem:

$$\begin{aligned} rV - \frac{\partial V}{\partial t} = \max_{z_M, z_L} \big\{ & \Pi_f - C_M(z_M) q_M - C_L(z_L) q_L \\ & + z_M \left[V(q_M(1 + \lambda_M), q_L) - V(q_M, q_L) \right] \\ & + z_L \left[V(q_M, q_L(1 + \lambda_L)) - V(q_M, q_L) \right] \big\} \end{aligned} \tag{33}$$

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 term $z_M \left[V(q_M(1 + \lambda_M), q_L) - V(q_M, q_L) \right]$ represents the net expected value from inno-

vation in automated tasks, while $z_L\left[V\left(q_M, q_L(1+\lambda_L)\right) - V\left(q_M, q_L\right)\right]$ represents the net expected value from innovation in non-automated tasks. The left-hand side of the equation accounts for the firm's value discounted at rate r.

We consider the stationary solution to the economy where $\frac{\partial V}{\partial t} = 0$ and the Markov Perfect Equilibrium, where the value function depends on the state variables q_M and q_L. Substituting the profit function into Equation (33), we obtain the Bellman equation:

$$\begin{aligned} rV = \max_{z_M, z_L} \Big\{ & \alpha_M \theta_f q_M + \alpha_L \left(1-\theta_f\right) q_L - C_M\left(z_M\right) q_M - C_L\left(z_L\right) q_L \\ & + z_M\left[V\left(q_M\left(1+\lambda_M\right), q_L\right) - V\left(q_M, q_L\right)\right] \\ & + z_L\left[V\left(q_M, q_L\left(1+\lambda_L\right)\right) - V\left(q_M, q_L\right)\right] \Big\} \end{aligned} \tag{34}$$

PROPOSITION 6.1. *The unique solution to Equation (34) is given by:*

$$V\left(q_M, q_L\right) = A_M q_M + A_L q_L, \tag{35}$$

where A_M and A_L are determined by:

$$\begin{cases} rA_M = \alpha_M \theta_f - C_M\left(z_M^*\right) + z_M^* A_M \lambda_M, \\ rA_L = \left(1-\theta_f\right) \alpha_L - C_L\left(z_L^*\right) + z_L^* A_L \lambda_L. \end{cases} \tag{36}$$

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

$$C_M'\left(z_M^*\right) = A_M \lambda_M, \quad \textit{and} \quad C_L'\left(z_L^*\right) = \lambda_L A_L. \tag{37}$$

6.4 Automation and Innovation Shift

The next proposition states how automation affects innovation in both automated and non-automated tasks. Specifically, automation across all production lines is denoted as:

$$\gamma = \theta_f \times \alpha_M,$$

where θ_f captures the extensive margin while α_M captures the intensive margin, as discussed in 6.1. We examine $\frac{\partial z_M^*}{\partial \gamma}$ and $\frac{\partial z_L^*}{\partial \gamma}$.

PROPOSITION 6.2. *A higher level of automation induces a higher level of innovation in automated tasks and a lower level of innovation in non-automated tasks, that is,*

$$\begin{aligned} \frac{\partial z_M^*}{\partial \gamma} &= \frac{\lambda_M}{rC_M''(z_M)} > 0, \\ \frac{\partial z_L^*}{\partial \gamma} &= -\frac{\lambda_L}{rC_L''(z_L)} \times \frac{\alpha_L}{\alpha_M} < 0. \end{aligned} \tag{38}$$

The following discuss several comparative analysis about empirical studies.

The effect of γ. Empirically, we interpret automation γ as the degree of robot exposure for firm f, which is driven by two key components: the proportion of tasks that are automated, represented by θ_f, and the intensity of automation within each automated task, captured by α_M in equilibrium.

To measure the arrival rate of innovations in automated tasks, z_M, we use patents related to AI, as they play a crucial role in facilitating the deployment of robots and enhancing their productivity. Similarly, we measure the arrival rate of innovations in non-automated tasks, z_L, using other patents.

The Effect of $C''(\cdot)$. For simplicity, we assume the cost function takes the quadratic form:

$$C_k(z) = \frac{c_k}{2} z^2, \quad k \in \{M, L\}.$$

Combining the results in Proposition 6.2, this specification suggests that firms with a lower c_M would exhibit greater sensitivity to rising robot exposure, shifting their innovation efforts more aggressively toward AI-driven innovations. If we interpret firms with a strong prior research focus on AI as having a lower c_M, we would expect these firms to reallocate their innovation efforts toward AI at a higher rate compared to firms with less AI-related experience.

The Effect of λ_M. The parameter λ_M can be interpreted as the quality of AI patent in terms of facilitating productivity. Firms with high-quality innovation - characterized by a higher λ_M

- will exhibit a stronger response to increased robot exposure by shifting their research efforts more significantly toward AI-related innovations.

7 Conclusions

Corporate innovations appear to have slowed down despite increases in R&D expenditure in recent years. We argue and show evidence that this is a manifestation of firms' endogenous responses to rising automation and robot exposure around the globe. As firms shift corporate innovations from incumbent fields towards AI-oriented ones that complement automation, they naturally witness a significant decline in patent output but a rise in the innovation inputs, an automation-induced innovation shift that is more significant for firms with more AI-related research experience or generate more data, and can be rationalized in a simple dynamic equilibrium model. Our study constitutes an initial investigation into the broader question of how one innovation or technology adoption induces other innovations in the short term and over the long term, which calls for future research.

References

Daron Acemoglu and Joshua Linn. Market size in innovation: theory and evidence from the pharmaceutical industry. *The Quarterly journal of economics*, 119(3):1049–1090, 2004.

Daron Acemoglu and Pascual Restrepo. The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6):1488–1542, 2018.

Daron Acemoglu and Pascual Restrepo. Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2):3–30, 2019.

Daron Acemoglu and Pascual Restrepo. Robots and jobs: Evidence from us labor markets. *Journal of Political Economy*, 128(6):2188–2244, 2020.

Daron Acemoglu and Pascual Restrepo. Demographics and automation. *The Review of Economic Studies*, 89(1):1–44, 2022.

Daron Acemoglu, Ufuk Akcigit, and William R Kerr. Innovation network. *Proceedings of the National Academy of Sciences*, 113(41):11483–11488, 2016.

Daron Acemoglu, Ufuk Akcigit, Harun Alp, Nicholas Bloom, and William Kerr. Innovation, reallocation, and growth. *American Economic Review*, 108(11):3450–91, 2018.

Viral V Acharya and Krishnamurthy V Subramanian. Bankruptcy codes and innovation. *The Review of Financial Studies*, 22(12):4949–4988, 2009.

Viral V Acharya, Ramin P Baghai, and Krishnamurthy V Subramanian. Wrongful discharge laws and innovation. *The Review of Financial Studies*, 27(1):301–346, 2014.

Philippe Aghion and Peter Howitt. A model of growth through creative destruction. *Econometrica (1986-1998)*, 60(2):323, 1992.

Ajay Agrawal, Joshua S Gans, and Avi Goldfarb. Artificial intelligence: the ambiguous labor market impact of automating prediction. *Journal of Economic Perspectives*, 33 (2):31–50, 2019.

Ufuk Akcigit and William R Kerr. Growth through heterogeneous innovations. *Journal of Political Economy*, 126(4):1374–1443, 2018.

Ashish Arora, Sharon Belenzon, and Andrea Patacconi. The decline of science in corporate r&d. *Strategic Management Journal*, 39(1):3–32, 2018.

David Autor and Anna Salomons. Is automation labor-displacing? productivity growth, employment, and the labor share. Technical report, National Bureau of Economic Research, 2018.

Tania Babina, Anastassia Fedyk, Alex He, and James Hodson. Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151:103745, 2024.

Patrick Bajari, Victor Chernozhukov, Ali Hortaçsu, and Junichi Suzuki. The impact of big data on firm performance: An empirical investigation. In *AEA papers and proceedings*, volume 109, pages 33–37. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, 2019.

Juliane Begenau, Maryam Farboodi, and Laura Veldkamp. Big data in finance and the growth of large firms. *Journal of Monetary Economics*, 97:71–87, 2018.

Hamid Bekamiri, Daniel S Hain, and Roman Jurowetzki. Patentsberta: A deep nlp based hybrid model for patent distance and classification using augmented sbert. *arXiv preprint arXiv:2103.11933*, 2021.

Jan Bena and Kai Li. Corporate innovations and mergers and acquisitions. *The Journal of Finance*, 69(5):1923–1960, 2014.

Martin Beraja, David Y Yang, and Noam Yuchtman. Data-intensive innovation and the state: Evidence from ai firms in china. *The Review of Economic Studies*, 2022.

Andrew Berg, Edward F Buffie, and Luis-Felipe Zanna. Should we fear the robot revolution?(the correct answer is yes). *Journal of Monetary Economics*, 97:117–148, 2018.

Nicholas Bloom, Mark Schankerman, and John Van Reenen. Identifying technology spillovers and product market rivalry. *Econometrica*, 81(4):1347–1393, 2013.

Nicholas Bloom, Charles I Jones, John Van Reenen, and Michael Webb. Are ideas getting harder to find? *American Economic Review*, 110(4):1104–1144, 2020.

Nicholas Bloom, Steven J Davis, and Yulia Zhestkova. Covid-19 shifted patent applications toward technologies that support working from home. In *AEA Papers and Proceedings*, volume 111, pages 263–66, 2021.

Kirill Borusyak, Peter Hull, and Xavier Jaravel. Quasi-experimental shift-share research designs. *The Review of economic studies*, 89(1):181–213, 2022.

Erik Brynjolfsson and Kristina McElheran. The rapid adoption of data-driven decision-making. *American Economic Review*, 106(5):133–139, 2016.

Erik Brynjolfsson and Tom Mitchell. What can machine learning do? workforce implications. *Science*, 358(6370):1530–1534, 2017.

Erik Brynjolfsson, L Hitt, and H Kim. How does data-driven decision making affect firm productivity and ceo pay. *Wharton, School, University of Pennsylvania*, 2010.

Francesco Caselli and Alan Manning. Robot arithmetic: new technology and wages. *American Economic Review: Insights*, 1(1):1–12, 2019.

Geraldo Cerqueiro, Deepak Hegde, María Fabiana Penas, and Robert C Seamans. Debtor rights, credit supply, and innovation. *Management Science*, 63(10):3311–3327, 2017.

Hsinchun Chen, Roger HL Chiang, and Veda C Storey. Business intelligence and analytics: From big data to big impact. *MIS quarterly*, pages 1165–1188, 2012.

Mark A Chen, Qinxi Wu, and Baozhong Yang. How valuable is fintech innovation? *The Review of Financial Studies*, 32(5):2062–2106, 2019.

Iain M Cockburn, Rebecca Henderson, and Scott Stern. The impact of artificial intelligence on innovation: An exploratory analysis. In *The economics of artificial intelligence: An agenda*, pages 115–146. University of Chicago Press, 2018.

Jonathan B Cohn, Zack Liu, and Malcolm I Wardlaw. Count (and count-like) data in finance. *Journal of Financial Economics*, 146(2):529–551, 2022.

Lin William Cong, Tengyuan Liang, Xiao Zhang, and Wu Zhu. Textual factors: A scalable, interpretable, and data-driven approach to analyzing unstructured information. *Interpretable, and Data-driven Approach to Analyzing Unstructured Information (September 1, 2019)*, 2019.

Arnaud Costinot, Dave Donaldson, Margaret Kyle, and Heidi Williams. The more we die, the more we sell? a simple test of the home-market effect. *The quarterly journal of economics*, 134(2):843–894, 2019.

Jay Dixon, Bryan Hong, and Lynn Wu. The robot revolution: Managerial and employment consequences for firms. *Management Science*, 67(9):5586–5605, 2021.

Florian Ederer and Gustavo Manso. Incentives for innovation: Bankruptcy, corporate governance, and compensation systems. In *Handbook on law, innovation and growth*. Edward Elgar Publishing, 2011.

Lily H Fang, Josh Lerner, and Chaopeng Wu. Intellectual property rights protection, ownership, and innovation: Evidence from china. *The Review of Financial Studies*, 30(7): 2446–2477, 2017.

Maryam Farboodi, Roxana Mihet, Thomas Philippon, and Laura Veldkamp. Big data and firm dynamics. In *AEA papers and proceedings*, volume 109, pages 38–42, 2019.

Amy Finkelstein. Static and dynamic effects of health policy: Evidence from the vaccine industry. *The Quarterly Journal of Economics*, 119(2):527–564, 2004.

Andreas Fuster, Matthew Plosser, Philipp Schnabl, and James Vickery. The role of technology in mortgage lending. *The Review of Financial Studies*, 32(5):1854–1899, 2019.

Alexander V Giczy, Nicholas A Pairolero, and Andrew A Toole. Identifying artificial intelligence (ai) invention: A novel ai patent dataset. *The Journal of Technology Transfer*,

47(2):476–505, 2022.

Itay Goldstein, Wei Jiang, and G Andrew Karolyi. To fintech and beyond. *The Review of Financial Studies*, 32(5):1647–1661, 2019.

Georg Graetz and Guy Michaels. Robots at work. *Review of Economics and Statistics*, 100(5):753–768, 2018.

Zvi Griliches. Productivity, r&d, and the data constraint. In *The Economic Impact of Knowledge*, pages 213–241. Routledge, 2009.

Joao Guerreiro, Sergio Rebelo, and Pedro Teles. Should robots be taxed? *The Review of Economic Studies*, 89(1):279–311, 2022.

Bronwyn H Hall, Adam B Jaffe, and Manuel Trajtenberg. The nber patent citation data file: Lessons, insights and methodological tools, 2001.

Jie He and Xuan Tian. Finance and corporate innovation: A survey. *Asia-Pacific Journal of Financial Studies*, 47(2):165–212, 2018.

Jie Jack He and Xuan Tian. The dark side of analyst coverage: The case of innovation. *Journal of Financial Economics*, 109(3):856–878, 2013.

Deepak Hegde, Alexander Ljungqvist, and Manav Raj. Quick or broad patents? evidence from us startups. *The Review of Financial Studies*, 35(6):2705–2742, 2022.

Po-Hsuan Hsu, Xuan Tian, and Yan Xu. Financial development and innovation: Cross-country evidence. *Journal of financial economics*, 112(1):116–135, 2014.

Anders Humlum. *Robot adoption and labor market dynamics*. Rockwool Foundation Research Unit Berlin, Germany, 2022.

Benjamin F Jones. The burden of knowledge and the "death of the renaissance man": Is innovation getting harder? *The Review of Economic Studies*, 76(1):283–317, 2009.

John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *nature*, 596(7873): 583–589, 2021.

Michael Koch, Ilya Manuylov, and Marcel Smolka. Robots and firms. *The Economic Journal*, 131(638):2553–2584, 2021.

Leonid Kogan, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. Technological innovation, resource allocation, and growth. *The quarterly journal of economics*, 132 (2):665–712, 2017.

Josh Lerner. Assessing the contribution of venture capital. *the RAND Journal of Economics*, 31(4):674–692, 2000.

Josh Lerner. The empirical impact of intellectual property rights on innovation: Puzzles and clues. *American Economic Review*, 99(2):343–48, 2009.

Chen Lin, Sibo Liu, and Gustavo Manso. Shareholder litigation and corporate innovation. *Management Science*, 67(6):3346–3367, 2021.

Song Ma. Technological obsolescence. 2023.

Ulrike Malmendier and Geoffrey Tate. Ceo overconfidence and corporate investment. *The journal of finance*, 60(6):2661–2700, 2005.

Andrew McAfee, Erik Brynjolfsson, Thomas H Davenport, DJ Patil, and Dominic Barton. Big data: the management revolution. *Harvard business review*, 90(10):60–68, 2012.

Filippo Mezzanotti and Timothy Simcoe. Research and/or development? financial frictions and innovation investment. Technical report, National Bureau of Economic Research, 2023.

Paul M Romer. Increasing returns and long-run growth. *Journal of political economy*, 94 (5):1002–1037, 1986.

Joseph Schumpeter and Ursula Backhaus. The theory of economic development. In *Joseph Alois Schumpeter*, pages 61–116. Springer, 2003.

James Stock and Motohiro Yogo. Asymptotic distributions of instrumental variables statistics with many instruments. *Identification and inference for econometric models: Essays in honor of Thomas Rothenberg*, 6:109–120, 2005.

Jayanthi Sunder, Shyam V Sunder, and Jingjing Zhang. Pilot ceos and corporate innovation. *Journal of Financial Economics*, 123(1):209–224, 2017.

Prasanna Tambe. Big data investment, skills, and firm value. *Management science*, 60(6): 1452–1469, 2014.

Lynn Wu, Lorin Hitt, and Bowen Lou. Data analytics, innovation, and firm productivity. *Management Science*, 66(5):2017–2039, 2020.

Yucheng Yang and Wu Zhu. Networks and business cycles. *Available at SSRN 3718826*, 2020.

Chris Yung. Making waves: To innovate or be a fast second? *Journal of Financial and Quantitative Analysis*, 51(2):415–433, 2016.

Joseph Zeira. Workers, machines, and economic growth. *The Quarterly Journal of Economics*, 113(4):1091–1117, 1998.

Miao Ben Zhang. Labor-technology substitution: Implications for asset pricing. *The Journal of Finance*, 74(4):1793–1839, 2019.

Christina Zhu. Big data as a governance mechanism. *The Review of Financial Studies*, 32

(5):2021–2061, 2019.

Figures and Tables

Figure 1: 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. We use the robot stock data from IFR and the employment data from BLS. And since the IFR only reports the robot applications in the U.S., Canada, and Mexico together at the industry level before 2011, we use the robots for these three countries following Acemoglu and Restrepo [2020], who documented that the U.S. accounts for more than 90 percent of the robot market share in North America.

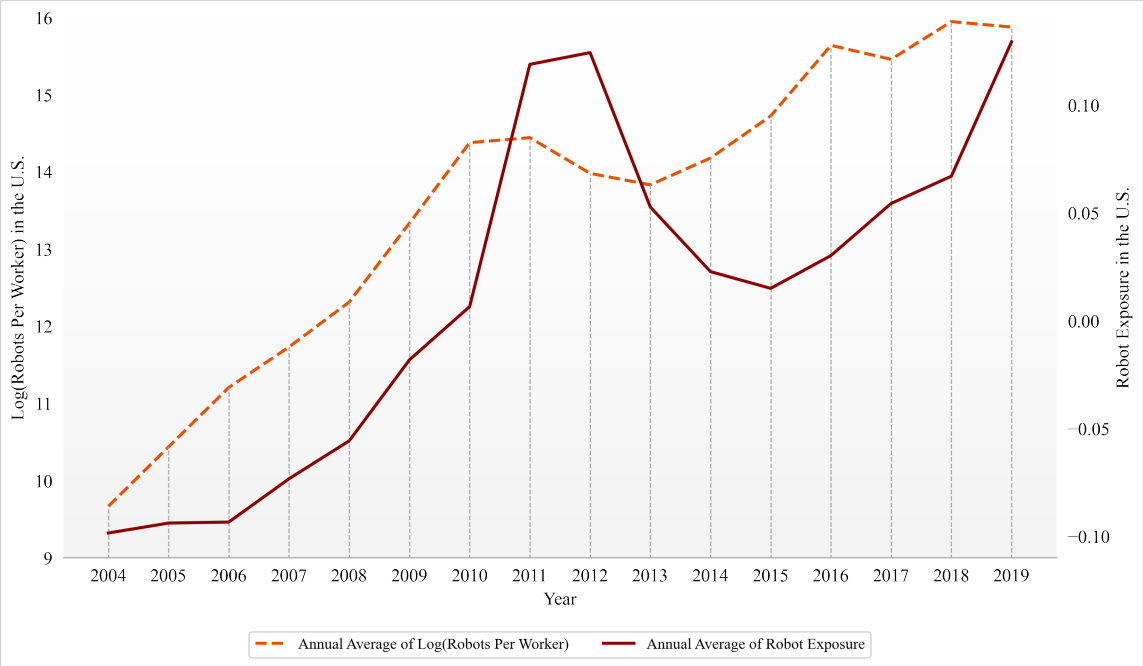
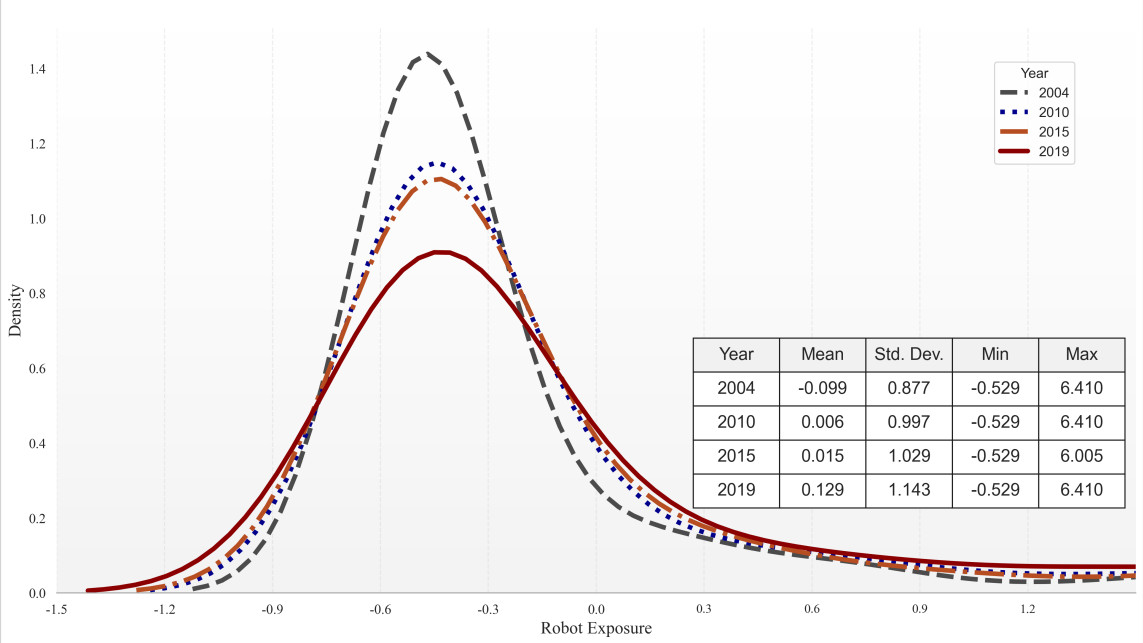


Figure 2: Distribution of Robot Exposure

This figure illustrates the distribution curves of firm-level robot exposure in the years 2004, 2010, 2015, and 2019. The 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 then standardized to have a mean of 0 and a variance of 1. Additionally, the figure presents the basic 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 3: Distribution of Incremental Changes in Robot Exposure for Top 10% Firms

This figure illustrates the incremental changes in firm-level 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 this change in robot exposure from the previous year. For these firms, we then compute the subsequent year's robot exposure change and present the distribution of these incremental changes.

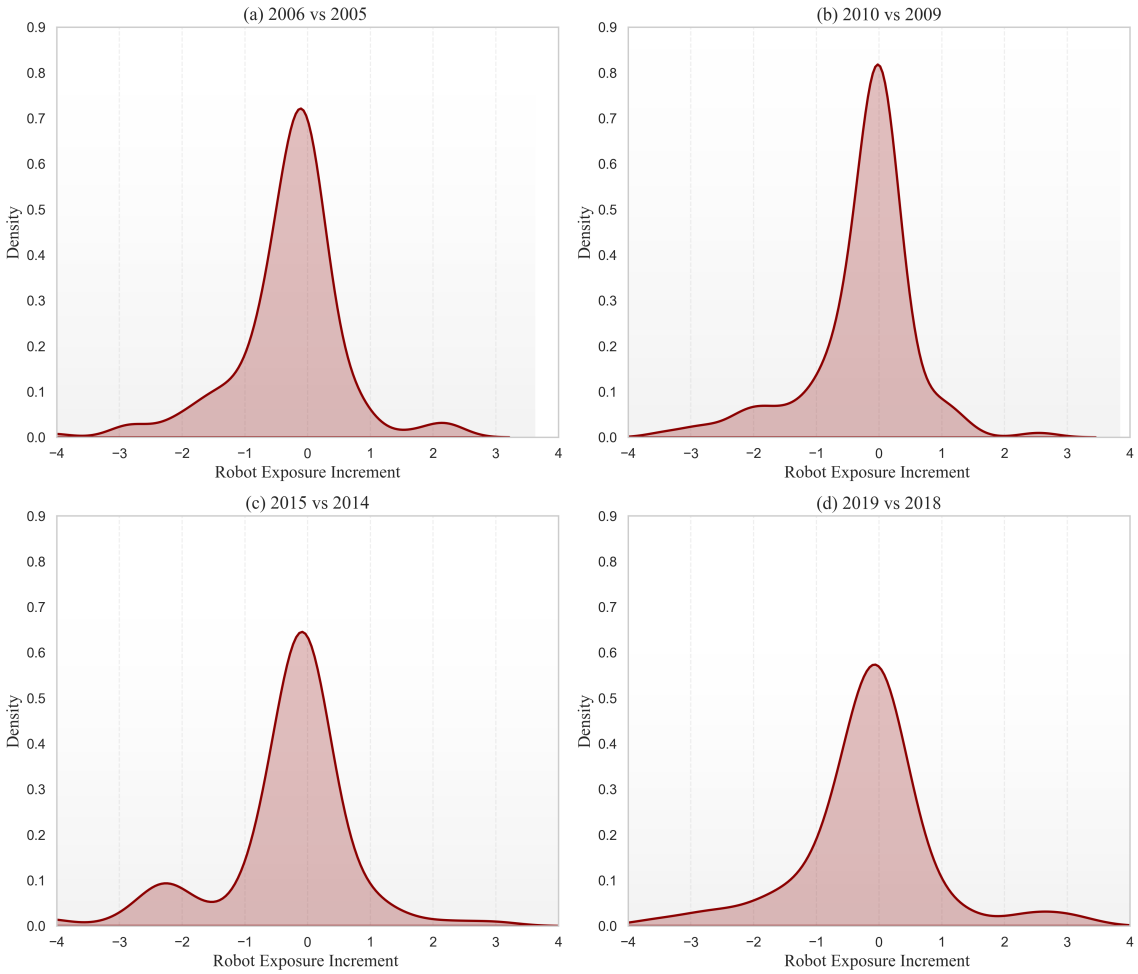


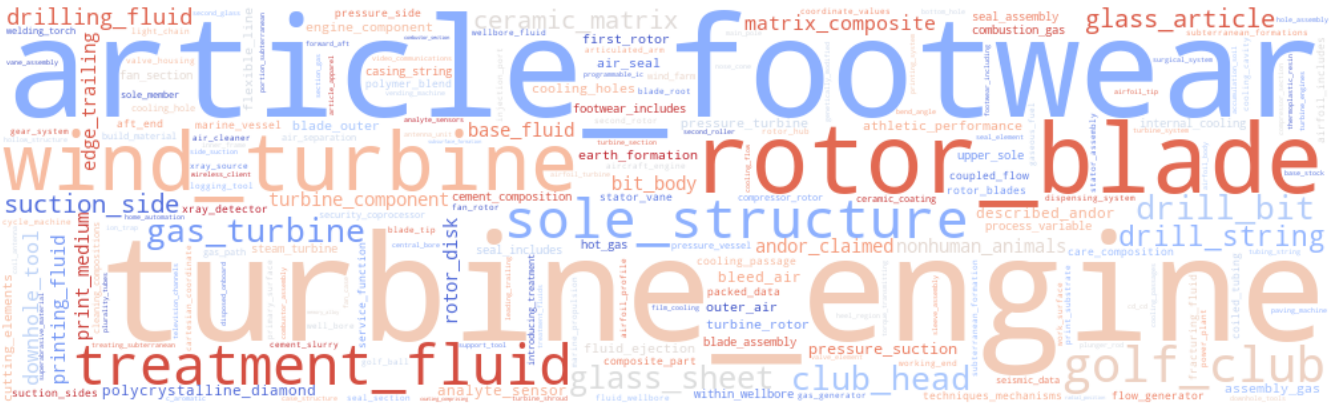
Figure 4: Word Cloud of Relative Frequency Differences between High and Low Robot Exposure Patents

This figure presents word clouds that display the relative frequency differences of n-grams between patents with robot exposure above the median and those below the median, based on samples from 2010, 2015, and 2019. Specifically, we first calculate the frequency values of the patent abstract pools for these two groups relative to the baseline year of 2004. Based on these frequencies, we compute the relative frequencies (after log-transformation) between the two groups, which are then visualized. The word cloud highlights n-grams with significant frequency differences in a given context. The size of the font in the word cloud represents the relative ratio of each n-gram, with larger words indicating n-grams with higher relative frequencies.

(a) 2010



(b) 2015



(c) 2019

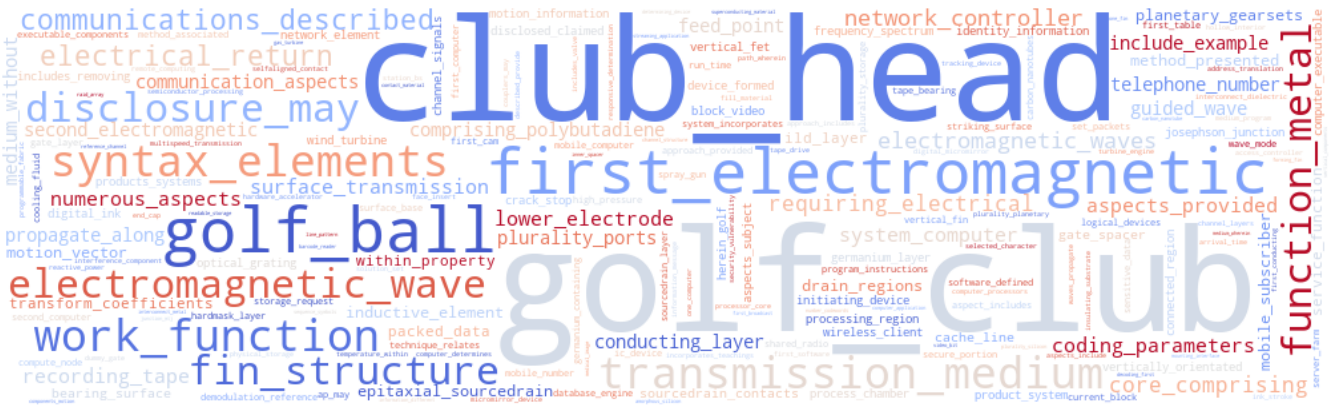


Figure 5: Keywords Frequencies of Patents

This figure reports the relative frequency of keywords for the filing years 2010, 2015, and 2019, using 2004 as the base year. To determine the relative frequency for each year (2010, 2015, and 2019), we first calculate the distribution of patents across keywords for that year and then normalize this distribution relative to the 2004 baseline.

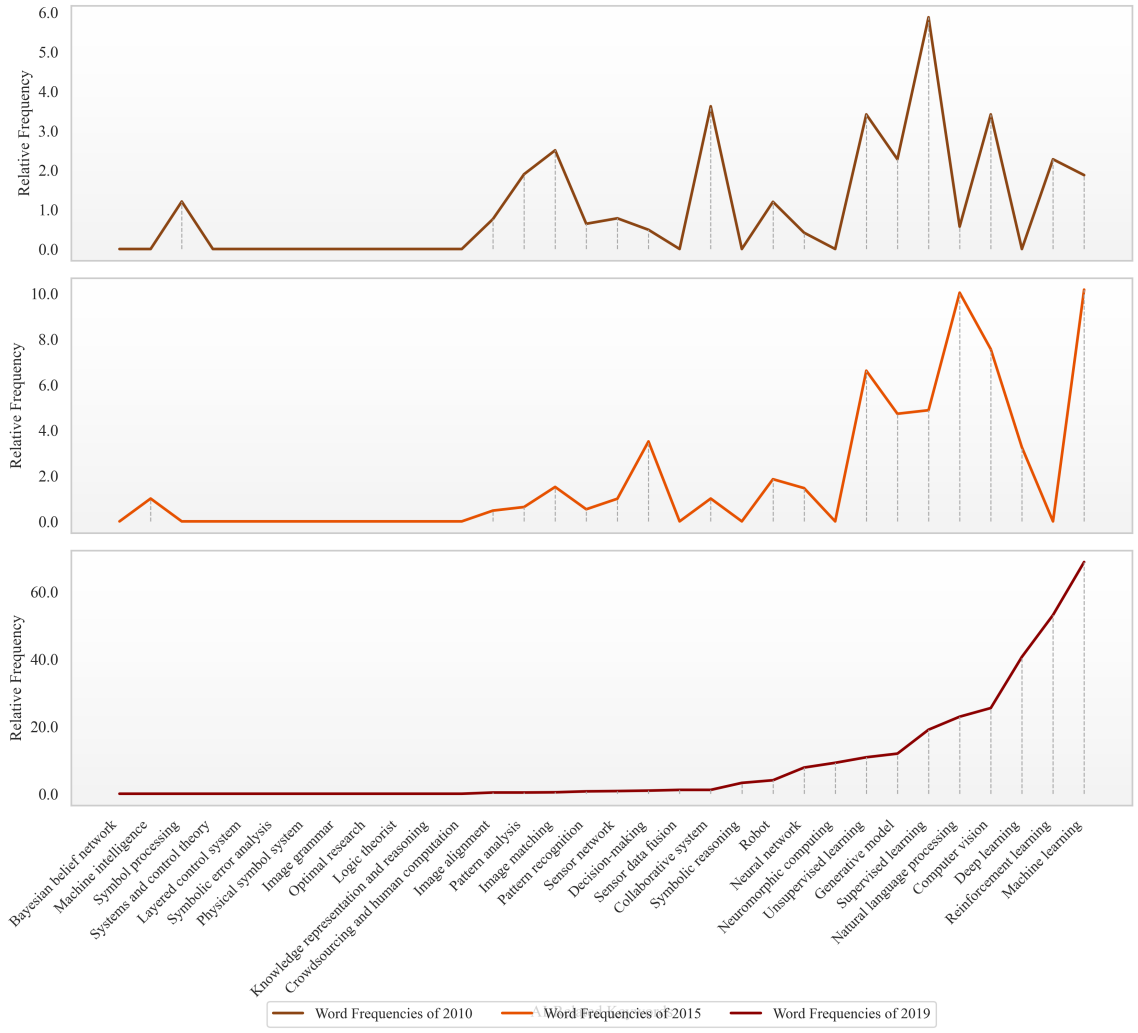


Figure 6: The Impact of Robot Exposure on Different Types of Patent Classes

This figure examines the impact of robot exposure on different patent classes. Specifically, we categorize patents by their two-digit CPC codes and introduce an interaction term between these categories and robot exposure in the following equation:

$$\mathrm{E}\left(\mathrm{y}_{\mathrm{it+j}} \mid \mathbf{X}_{\mathrm{it}}\right) = \exp\left(\text{Robot Exposure}_{it}\beta + \mathbf{X}'_{it}\gamma\right)$$

The figure illustrates the top 10 CPC industries with significantly negative coefficients for the interaction term, as well as the corresponding interaction term coefficients for the industries with significantly positive coefficients.

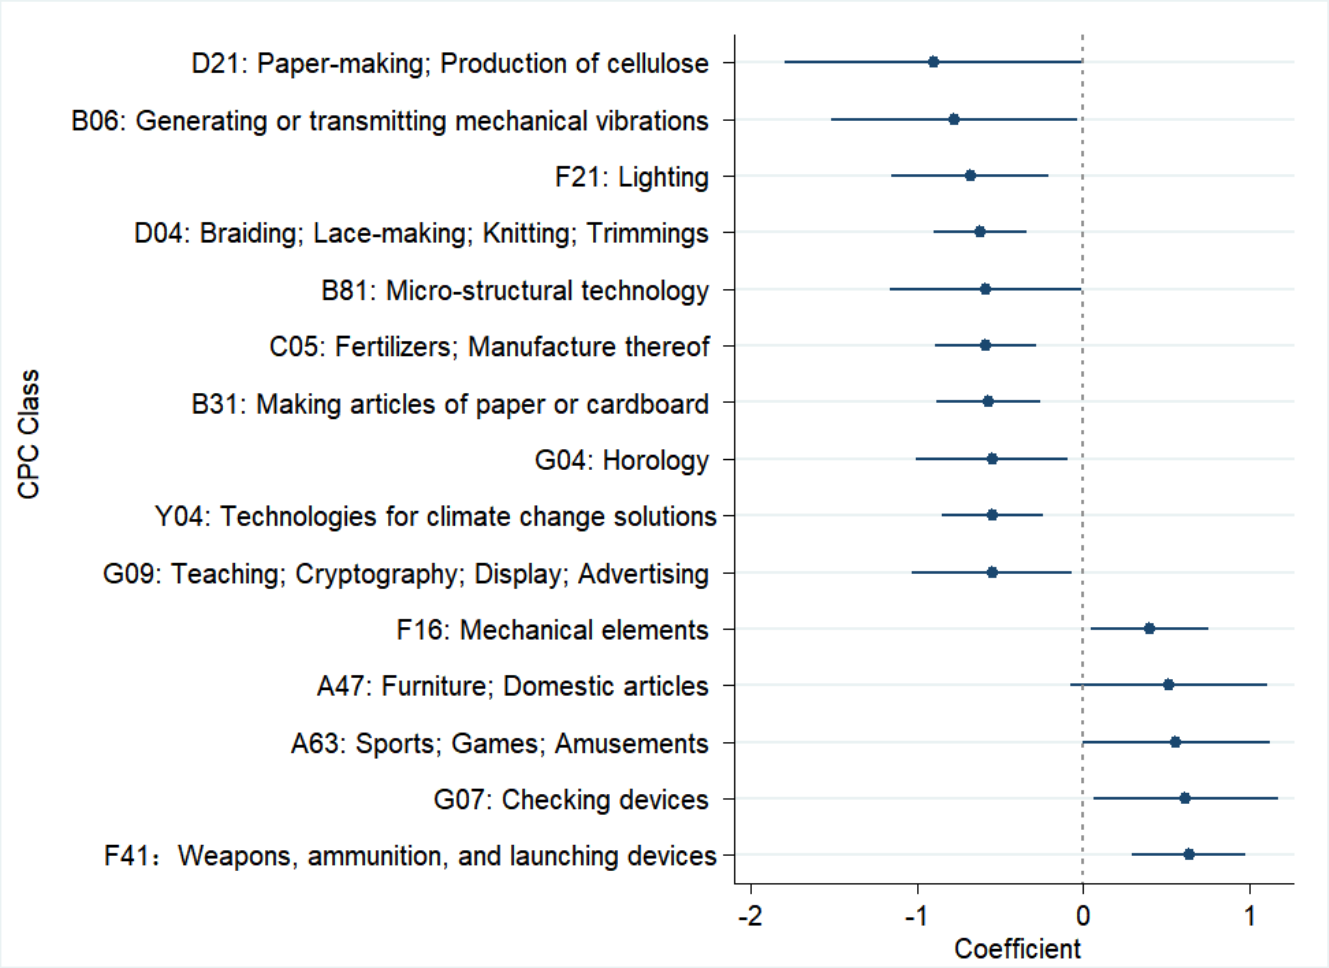


Figure 7: Comparison of AI and Non-AI Patents

This table presents a comparison of AI and non-AI patents across four dimensions: team size, labor input, inventor originality, and patent value, between the filing years 2004 and 2019. The calculation of team size for granted patents follows Jones [2009], which defines team size as the number of inventors involved in the patent. For labor input, if an inventor applies for two patents in the same year, the labor input for each patent is adjusted to 0.5. The labor input is then aggregated at the patent level. Inventor originality for granted patents is calculated following Hall et al. [2001], based on the method outlined below. We measure an inventor's originality as $Inventor\ Originality_{i,t} = 1 - \sum_{j=1}^{N} s_{ji,t}^2$, where $s_{ji,t}$ is the fraction of patents by inventor i that fall into technology j over the past five years. Patent value is measured following the approach proposed by Kogan et al. [2017]. In the figure, the solid crimson line represents the temporal trend of the annual mean for AI patents, while the gray dashed line represents the temporal trend for non-AI patents.

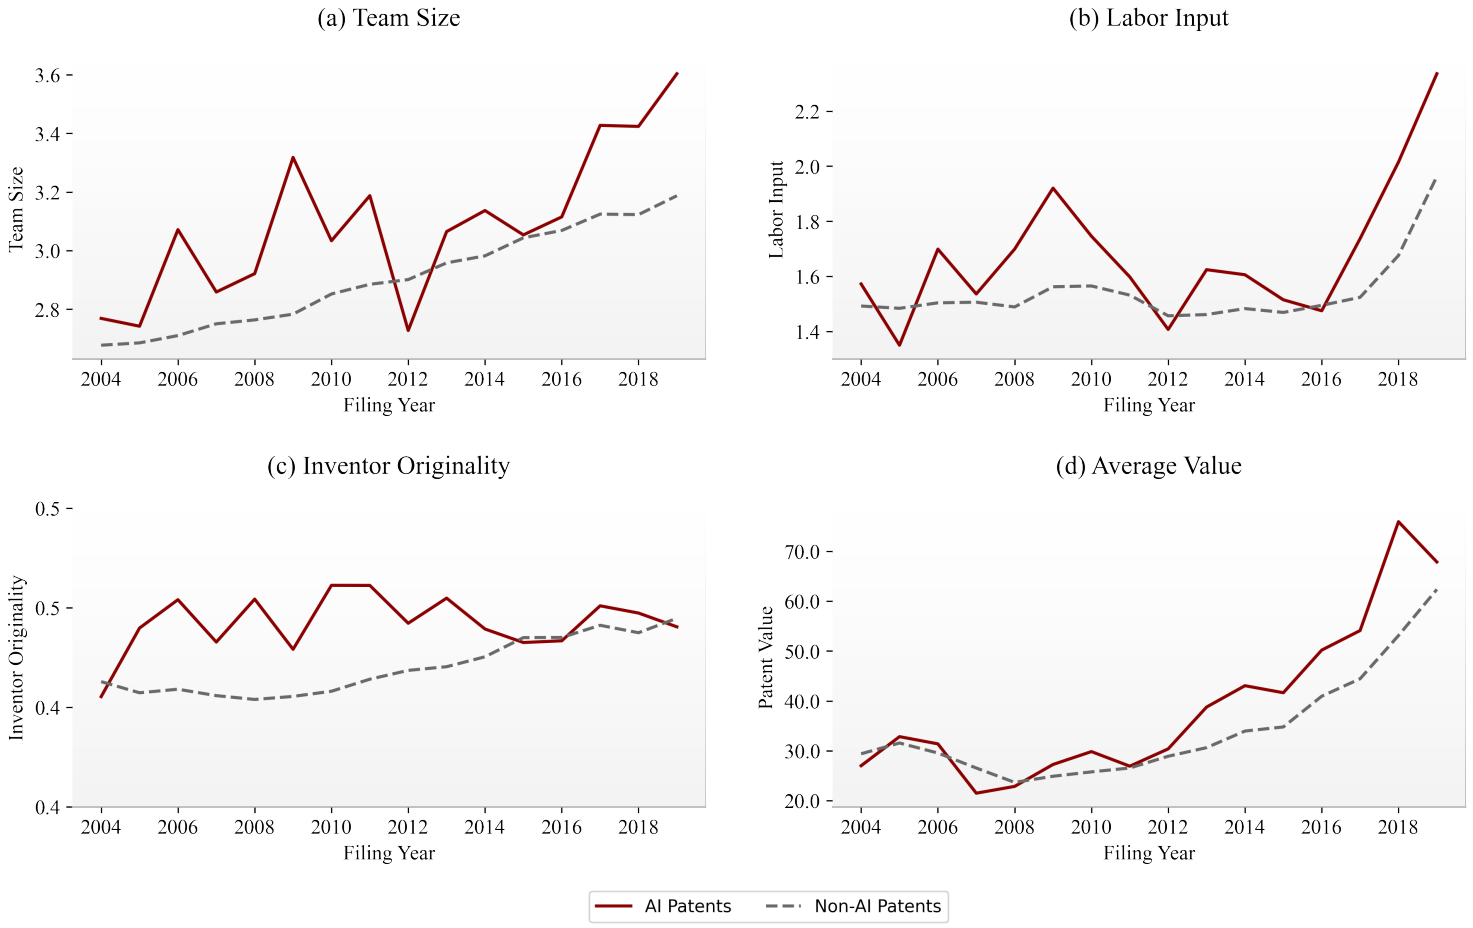


Table 1: Summary of the Match between CompuStat and USPTO

This table provides a detailed summary of the matching process between Compustat and USPTO data. Column (1) lists the years, while Column (2) reports the annual count of Compustat firms that hold at least one patent throughout the sample period. Columns (3) and (4) present the yearly count and corresponding percentage of firms with at least one patent, respectively. Finally, Column (5) displays the total number of patents granted to Compustat firms each year.

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 2: Keyword List for Defining Robot Exposure

This table presents the keyword list used to define firm-level robot exposure. The list is constructed based on a comprehensive review of relevant literature on robot adoption, with a substantial portion derived from IFR Reports, which offer detailed definitions of industrial robots and their applications. Column (1) cites the sources, while the remaining columns display the keywords. When counting keyword frequencies, we do not differentiate by case sensitivity or singular and plural forms.

References	Keywords		
IFR Report	Adopt robot	Industrial robot	Robot-driven
Acemoglu and Restrepo (2020)	Advanced robot	Innovative robot	Robot employ
Acemoglu and Restrepo (2022)	AI automation	Install robot	Robot enabled
Zeira (1998)	AI robot	Intelligent robot	Robot exposure
Guerreiro et al. (2022)	AI-enabled robot	Interventional robot	Robot gripper
Dixon et al. (2021)	Articulated robot	Labor-replacing machine	Robot head
Graetz et al. (2018)	Artificial Intelligence	Large-scale robot	Robotic arm
Acemoglu and Restrepo (2018)	Assemble robot	Linear robot	Robotic hardware
Acemoglu and Restrepo (2019)	Assignment robot	Logistics robot	Robotic platform
Caselli and Manning (2019)	Automation application	Machine tending	Robotic software
Berg et al. (2018)	Automatically controlled manipulator	Maintain robot	Robotics related
Michael et al. (2021)	Automation-led	Manufacture robot	Robot industry
	Automation tech	More robot	Robot innovation
	Automative machine	Multipurpose manipulator	Robot input
	Automative robot	Multiple robot	Robot install
	Automative task	Multi-robot	Robot intensity
	Automation production	On-robot	Robot introduction
	Automative robot	Polar robot	Robot invest
	Automative system	Process robot	Robot labor
	Automative technology	Program robot	Robot manufaction
	Collaborative robot	Reprogrammable manipulator	Robot module
	Collaborative system	Robot innovation	Robot motor
	Computer-assisted machine	Rise of robot	Robot performance
	Co-robot	Robot adopt	Robot price
	Create robot	Robot advance	Robot-related
	Cylindrical robot	Robot application	Robot revolution
	Delta robot	Robot arithmetic	Robot-specialized
	Develop robot	Robot arm	Robot stock
	Disassemble robot	Robot as a service	Robot supply
	Dismantling robot	Robot assisted	Robot supporting
	Employ robot	Robot at work	Robot system
	Exposure to robot	Robot automation	Robot technology

Table 3: 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:

$$Technology\ Similarity_{i,t+j} = \alpha_i + \alpha_t + \beta \times Robot\ Exposure_{it} + X'_{it}\gamma + \varepsilon_{i,t},$$

where $Technology\ Similarity_{i,t+j}$ is measured using both $Tech\ Similarity_{i,t+j}$ and $BERT\ Similarity_{i,t+j}$. $Tech\ Similarity_{i,t+j}$ is quantified by the cosine similarity of the firm i's patent distributions between years $t+j$ and $t+j+5$. $BERT\ Similarity_{i,t+j}$ is calculated using the cosine similarity of the firm i's 768-dimensional continuous vector generated by the PatentBERT model between years $t+j$ and $t+j+5$. $Robot\ Exposure_{it}$ 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. Column (1) presents the first-stage result. Columns (2)–(4) report the second-stage estimates of the IV results for the coefficients of robot exposure on *Tech Similarity* for years t, $t+1$, and $t+2$, respectively. Columns (5)–(7) display the second-stage estimates of the IV results for the coefficients of robot exposure on *BERT Similarity* for years t, $t+1$, and $t+2$, 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 is at the firm-year level. All regressions control for firm-fixed effects and year-fixed effects. Robust standard errors are reported in parentheses. Coefficients marked with *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Level	Firm-Year						
Model	IV						
	First Stage	Second Stage					
Variable	Robot Exposure	Tech Similarity			BERT Similarity		
Period	t	t	t+1	t+2	t	t+1	t+2
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Robot Exposure (IV)	0.004***						
	(0.001)						
Robot Exposure		-0.187**	-0.213*	-0.263*	-0.460**	-0.722**	-0.812*
		(0.093)	(0.123)	(0.158)	(0.229)	(0.363)	(0.466)
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	20,815	20,815	18,088	15,908	17,915	15,625	13,793
Adjusted R-Squared	0.0583						
Week Identification Test F Statistics	24.00						
Underidentification Test P-value	0.000						
Endogeneity Test P-value	0.001						

Table 4: Keyword List for Defining AI Patents

This table lists the AI-related keywords used to identify AI patents, encompassing three main categories: robotic, symbolic system, and neural network. For details on the selection process, refer to Cockburn et al. (2019). AI patents are defined as those containing AI keywords or those directly citing patents that include AI keywords. In the analysis, we do not distinguish between case sensitivity or singular and plural forms of these keywords.

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

Table 5: The Impact of Robot Exposure on AI-Directed Innovation Shift

This table presents an analysis of whether firms adjust their innovation activities toward AI-related areas in response to increased robot exposure. Specifically, Columns (1)–(6) report the second-stage results of the 2SLS estimation, where we estimate the following equation:

$$y_{i,p,t+j} = \alpha_i + \alpha_t + \beta \times Robot\ Exposure_{it} + X'_{it}\gamma + \varepsilon_{i,p,t},$$

where $y_{i,p,t+j}$ is measured using both $AI\ Distance_{i,p,t+j}$ and $AI\ Closeness_{i,p,t+j}$. For each AI field, $AI\ Distance_{i,p,t+j}$ represents the minimum number of steps required for firm i's patent p to reach AI patents within the citation network in year $t+j$. $AI\ Closeness_{i,p,t+j}$ is calculated as the cosine similarity between the citation distribution vectors of firm i's patent p and AI patents in year $t+j$. $Robot\ Exposure_{it}$ 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. In Columns (7)–(9), we model the AI patent grant as a time-varying Poisson process and estimate the underlying arrival rate as:

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

where $\mathrm{E}\left(\mathrm{y_{i,t+j}} \mid \mathbf{X_{it}}\right)$ represents the expected AI patent numbers filed by firm i in year $t+j$. AI patents are identified as those containing AI-related keywords or those directly citing patents that include AI keywords listed in Table 4. Columns (1)–(3) report the coefficients of robot exposure on *AI Distance* for years t, $t+1$, and $t+2$, respectively. Columns (4)–(6) present the coefficients on *AI Closeness*, and Columns (7)–(9) display the results for *AI Patent Number* over the same time periods. 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 used in Columns (1)–(6) is at the patent-firm-year level, while the data in Columns (7)–(9) is at the firm-year level. All regressions incorporate firm-fixed and year-fixed effects. Robust standard errors are provided in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Level	Patent-Firm-Year						Firm-Year		
Model	IV						Possion		
Variable	AI Distance			AI Closeness			AI Patent Number		
Period	t	t+1	t+2	t	t+1	t+2	t	t+1	t+2
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Robot Exposure	-1.004***	-0.581***	-0.556***	0.079***	0.084***	0.087***	0.211***	0.244***	0.170**
	(0.124)	(0.124)	(0.140)	(0.004)	(0.004)	(0.004)	(0.066)	(0.077)	(0.085)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	661,018	607,772	561,545	661,018	607,772	561,545	3,058	2,692	2,396

Table 6: 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. Here, we consider 8 AI fields, including machine learning model (ML); evolutionary computation model (EVO); natural lang processing model (NLP); speech model (SPEECH); vision model (VISION); knowledge processing model (KR); planning/control model (PLAN); AI hardware model (HW). We devide each AI patents into one or more of these 8 fields and calculate the distance/closeness for any given patents to the specific filed AI patents. Specifically, we estimate:

$$y_{i,p,t} = \alpha_i + \alpha_t + \beta \times Robot\ Exposure_{it} + X'_{it}\gamma + \varepsilon_{i,p,t},$$

where $y_{i,p,t}$ is measured using both $AI\ Distance_{i,p,t}$ and $AI\ Closeness_{i,p,t}$. For each AI field, $AI\ Distance_{i,p,t}$ represents the minimum number of steps required for firm i's patent p to reach AI patents within the citation network in year t. $AI\ Closeness_{i,p,t}$ is calculated as the cosine similarity between the citation distribution vectors of firm i's patent p and AI patents in year t. $Robot\ Exposure_{it}$ 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. The dependent variables in Panel A and Panel B are *AI Closeness* and *AI Distance*, respectively. In each panel, Columns (1)–(8) present the IV coefficients of *Robot Exposure* in year t on eight types of AI-related measures for 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 sample period spans 2004–2019, and the regression data is at the patent-firm-year level. All regressions incorporate firm-fixed and year-fixed effects. Robust standard errors are provided in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Level	Patent-Firm-Year							
Model	IV							
Period	t							
Types	ML	EVO	NLP	SPEECH	VISION	KR	PLAN	HW
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Panel A: AI Closeness							
Variable	AI Closeness							
Robot Exposure	0.098***	0.060***	0.145***	0.136***	0.095***	0.145***	0.069***	0.122***
	(0.003)	(0.002)	(0.004)	(0.004)	(0.003)	(0.004)	(0.002)	(0.003)
Observations	607,772	607,772	607,772	607,772	607,772	607,772	607,772	607,772
	Panel B: AI Distance							
Variable	AI Distance							
Robot Exposure	-0.002	-0.018*	-0.049**	-0.019	-0.006	-0.015	-0.025	-0.447*
	(0.028)	(0.010)	(0.020)	(0.030)	(0.029)	(0.031)	(0.023)	(0.240)
Observations	607,772	607,772	607,772	607,772	607,772	607,772	607,772	607,772
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 7: The Impact of Robot Exposure on Data Generation

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

$$Data\ Generation_{i,t+j} = \alpha_i + \alpha_t + \beta \times Robot\ Exposure_{it} + X'_{it}\gamma + \varepsilon_{i,t},$$

where $Data\ Generation_{i,t+j}$ represents the frequency of data-generation-related keywords in firm i's 10-K report for year $t+j$, 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)–(3) present the coefficients of robot exposure on *Data Generation* for years t, $t+1$, and $t+2$, respectively. Columns (4) and (5) present the results of the heterogeneity analysis for large firms. The independent variable is the interaction between the *Dummy* variable and the standardized *Robot Exposure*. The *Dummy* variable indicates whether a firm qualifies as a large firm. In Column (4), 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. In Column (5), large firms are defined using *Dummy* $(Net\ Income > 75\%)$, which equals one if the firm's net profit 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 is at the firm-year level. All regressions control for firm-fixed effects and year-fixed effects. Robust 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				
Variable	Data Generation				
Period	t	t+1	t+2	t	t
	(1)	(2)	(3)	(4)	(5)
Robot Exposure	1.313***	1.428**	1.533**	1.043**	1.008**
	(0.485)	(0.650)	(0.772)	(0.476)	(0.486)
Robot Exposure × Dummy (Sales > 75%)				0.739*	
				(0.439)	
Robot Exposure × Dummy (Net Income > 75%)					0.587**
					(0.289)
Dummy (Sales > 75%)				-0.032	
				(0.044)	
Dummy (Net Income > 75%)					-0.049*
					(0.029)
Controls	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	20,815	18,090	15,909	20,815	20,815

Table 8: The Impact of Robot Exposure on Innovation Outputs

This table examines the impact of robot exposure on firms' innovation outputs in subsequent years. We consider two measures of innovation outputs: patent number and patent value. As for the patent number, we model the patent grant as a time-varying Poisson process and estimate the underlying arrival rate as:

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

where $\mathrm{E}\left(\mathrm{y_{i,t+j}} \mid \mathbf{X_{it}}\right)$ represents the expected patent numbers filed by firm i in year $t+j$. Here, *Citation-Deflated Patent Number* follows the method proposed by Acemoglu et al. [2016], which deflates the number of citations each patent received over the subsequent five years, normalizing these citations by the average within each 3-digit CPC class. $\textit{Robot Exposure}_{it}$ 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. As for the patent value, we estimate the following equation:

$$\textit{Citation-Deflated Patent Value}_{i,t+j} = \alpha_i + \alpha_t + \beta \times \textit{Robot Exposure}_{it} + X'_{it}\gamma + \varepsilon_{i,t},$$

where patent value is measured following the methodology proposed by Kogan et al. [2017]. We then deflate the number of citations each patent receives over the subsequent five years, normalize them by the average within each 3-digit CPC class, and aggregate the resulting patent values at the firm level. Columns (1)–(3) report the coefficients of robot exposure on *Citation-Deflated Patent Number* for years t, $t+1$, and $t+2$, respectively. Columns (4)–(6) display the the second-stage coefficients of robot exposure on *Citation-Deflated Patent Value* for years t, $t+1$, and $t+2$, 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 is at the firm-year level. All regressions control for firm-fixed effects and year-fixed effects. Error terms are adjusted for heteroskedasticity and clustered at the firm level. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Level	Firm-Year					
Model	Possion			IV		
Variable	Citation-Deflated Patent Number			Citation-Deflated Patent Value		
Period	t	t+1	t+2	t	t+1	t+2
	(1)	(2)	(3)	(4)	(5)	(6)
Robot Exposure	-0.086**	-0.041	-0.027	-0.997**	-1.165*	-1.601*
	(0.035)	(0.038)	(0.036)	(0.475)	(0.659)	(0.901)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	16,473	14,161	12,269	20,815	18,088	15,908

Table 9: The impact of Robot Exposure on the Cumulative Innovation Outputs

This table analyzes the effect of robot exposure on cumulative innovation outputs over multiple years. Panel A and Panel B display the results for the number of patents and patent value, respectively. Within each panel, Columns (1)–(6) report the Poisson and 2SLS regression estimates for the cumulative number of patents and cumulative patent value over periods ranging from 1 to 6 years, respectively. We model the patent grant as a time-varying Poisson process and estimate the underlying arrival rate as:

$$\mathrm{E}\left(\mathrm{y}_{\mathrm{i,t+j}} \mid \mathbf{X}_{\mathrm{it}}\right) = \exp\left(\beta_1 \times \textit{Robot Exposure}_{it} + \mathbf{X}'_{it}\gamma\right),$$

where $\mathrm{E}\left(\mathrm{y}_{\mathrm{i,t+j}} \mid \mathbf{X}_{\mathrm{it}}\right)$ represents the expected patent numbers filed by firm i in year $t+j$. Here, *Citation-Deflated Patent Number* follows the method proposed by Acemoglu et al. [2016], which deflates the number of citations each patent received over the subsequent five years, normalizing these citations by the average within each 3-digit CPC class. $\textit{Robot Exposure}_{it}$ 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. As for the patent value, we estimate the following equation:

$$\textit{Citation-Deflated Patent Value}_{i,t+j} = \alpha_i + \alpha_t + \beta \times \textit{Robot Exposure}_{it} + X'_{it}\gamma + \varepsilon_{i,t},$$

where patent value is measured following the methodology proposed by Kogan et al. [2017]. We then deflate the number of citations each patent receives over the subsequent five years, normalize them by the average within each 3-digit CPC class, 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 is at the firm-year level. All regressions control for firm-fixed effects and year-fixed effects. Error terms are adjusted for heteroskedasticity and clustered at the firm level. 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	Possion					
Variable	***Panel A: Citation-Deflated Patent Number***					
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
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table 10: The Impact of Robot Exposure on R&D Expenditure

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

$$y_{i,t+j} = a_i + a_t + \beta \times Robot\ Exposure_{i,t} + X'_{it}\gamma + \varepsilon_{i,t},$$

where $y_{i,t+j}$ represents firm i's R&D expenditure ratio in year $t+j$. 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. Columns (1)–(3) present the coefficients of robot exposure on $R\&D/Total\ Assets$ for years t, $t+1$, and $t+2$, respectively. Columns (4)–(6) show the coefficients of robot exposure on $R\&D/Sales$ for years t, $t+1$, and $t+2$, 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 is at the firm-year level. All regressions control for firm-fixed effects and year-fixed effects. Robust 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					
Variable	R&D/Total Assets			R&D/Sales		
Period	t	t+1	t+2	t	t+1	t+2
	(1)	(2)	(3)	(4)	(5)	(6)
Robot Exposure	0.192**	0.207**	0.182*	0.626***	0.684**	0.784**
	(0.080)	(0.102)	(0.109)	(0.230)	(0.296)	(0.382)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	20,815	18,088	15,908	20,815	18,088	15,908

Table 11: The Effect of Prior AI Research Experience on AI-Directed Innovation Shift

This table presents the second-stage results of the 2SLS estimation, which investigates the relationship between firms' prior AI research experience and AI-directed innovation shift. Specifically, we estimate the following equation:

$$y_{i,p,t} = \alpha_i + \alpha_t + \beta_0 \times \textit{Robot exposure}_{it} + \beta_1 \times \textit{AI experience}_{i,03} \times \textit{Robot exposure}_{it} + X'_{it}\gamma + \varepsilon_{i,p,t},$$

where $y_{i,p,t}$ is the AI distance (closeness) of firm i's patent p at year t. $\textit{AI Distance}_{i,p,t}$ represents the minimum number of steps required for firm i's patent p to reach AI patents within the citation network in year t. $\textit{AI Closeness}_{i,p,t}$ is calculated as the cosine similarity between the citation distribution vectors of firm i's patent p and AI patents in year t. $\textit{Robot Exposure}_{it}$ measures the degree of robot exposure by firm i in year t. $\textit{AI experience}_{i,03}$ 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. Column (1) and (2) present the coefficients on *AI Distance* and *AI Closeness* 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 is at the patent-firm-year level. All regressions control for firm-fixed effects and year-fixed effects. Robust 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	IV	
Period	t	
Variable	AI Distance	AI Closeness
	(1)	(2)
Robot Exposure × Dummy (AI Patent Before 2004 > Median)	-0.544***	0.053***
	(0.190)	(0.005)
Robot Exposure	-0.574***	0.050***
	(0.150)	(0.004)
Controls	Yes	Yes
Firm FE	Yes	Yes
Year FE	Yes	Yes
Observations	637,672	637,672

Table 12: The Impact of Robot Exposure on Patent Generality/Originality

This table presents the second-stage results of the 2SLS estimation, examining how a firm's robot exposure affects its patent generality and originality. Specifically, we estimate:

$$y_{i,p,t+j} = \alpha_i + \alpha_t + \beta \times Robot\ Exposure_{it} + X'_{it}\gamma + \varepsilon_{i,p,t},$$

where $y_{i,p,t+j}$ represents the innovation quality of firm i's patent p in year $t+j$. *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. Columns (1)–(3) report the coefficients of robot exposure on *Generality* for years t, $t+1$, and $t+2$, respectively. Columns (4)–(6) display the coefficients of robot exposure on *Originality* for years t, $t+1$, and $t+2$, 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, and the regression data is at the patent-firm-year level. All regressions incorporate firm-fixed and year-fixed effects. Robust standard errors are provided in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Level	Patent-Firm-Year					
Model	IV					
Variable	Generality			Originality		
Period	t	t+1	t+2	t	t+1	t+2
	(1)	(2)	(3)	(4)	(5)	(6)
Robot Exposure	0.040***	0.022***	0.022***	0.036***	0.018***	0.018***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	607,979	556,885	556,885	365,009	330,468	330,468

Appendix A: Figures

Figure A1: AI Patents in the US from 2004 to 2019

This figure reports AI patents as a percentage of total U.S. patents from filing years 2004 to 2019. AI patents are defined as those containing AI-related keywords or directly citing patents with AI keywords.

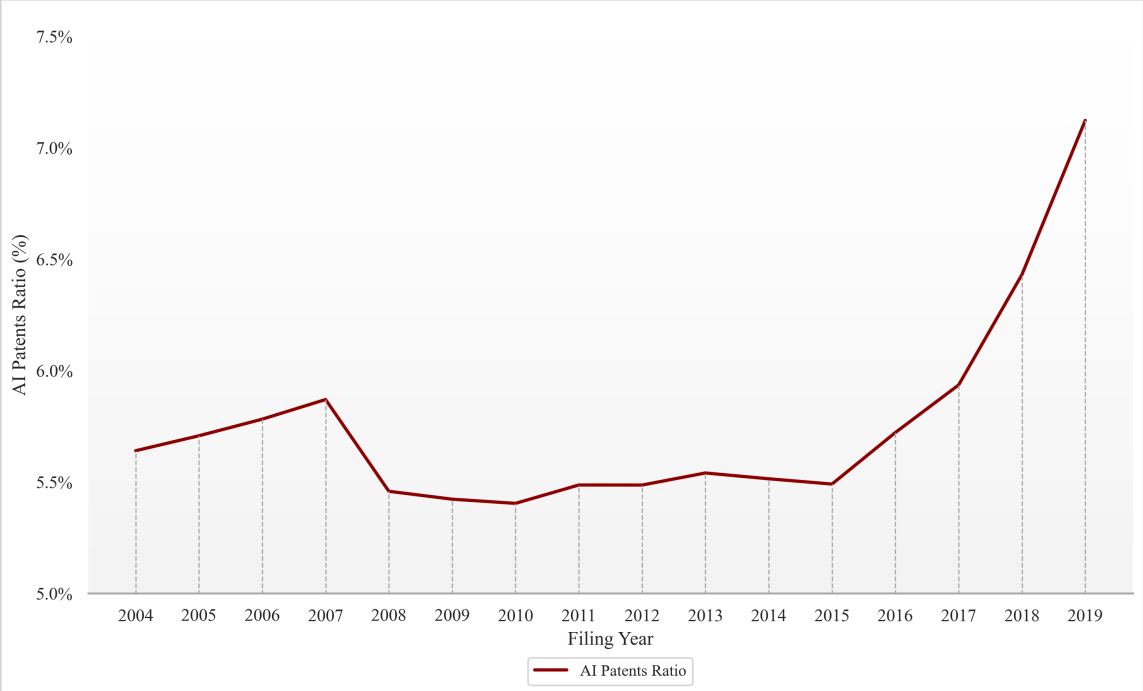
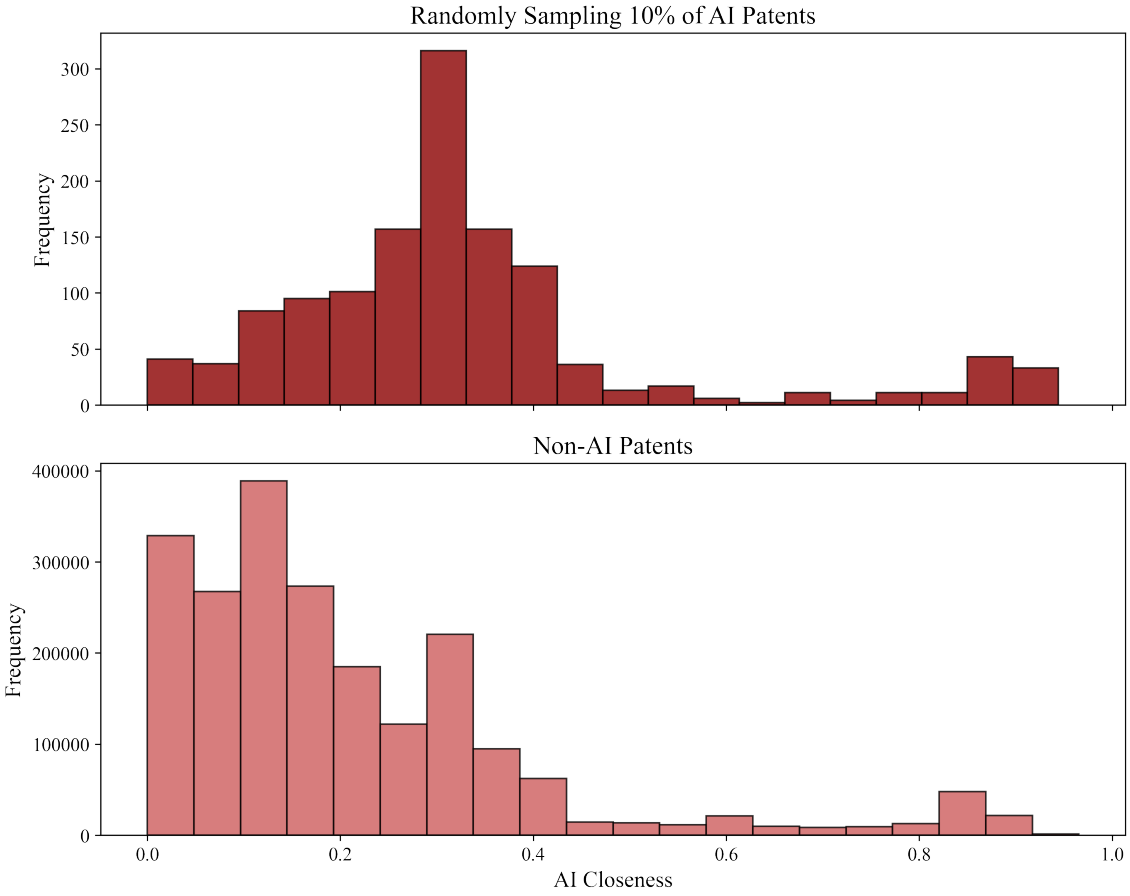


Figure A2: Difference in AI Closeness between AI Patents and Non-AI Patents

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 figure presents the resulting AI closeness metrics in histograms for comparison between the randomly selected AI patents and non-AI patents.



Appendix B: Tables

Table A1: Variable Definitions

Variable	Definition
Panel A: Firm Level	
Robot Variable	
Robot Exposure	We systematically reviewed all 10-K reports from 2004 to 2019, identifying mentions of robot-related keywords that directly pertain to the application or integration of robotic technologies within firms. Then we analyzed the frequency of these keywords in 10-K filings and normalized this frequency by the total word count in each report to account for the variation in document length: $$\text{Robot Exposure}_{i,t} = \sum_{w} \frac{\text{Robot Word Count}_{w,i,t}}{\text{Total Word Count}_{i,t}}$$ where $\text{Robot Word Count}_{w,i,t}$ represents the occurrences of keyword w for firm i in year t, and $\text{Total Word Count}_{i,t}$ is the total number of words in the same report.

Variable	Definition
Robot Exposure (IV)	First, we measure the robots using the data from the IFR. Then, combining the robot and employment in each year, we construct the instrumental variable following: $$Robot\ Exposure\ IV_{i,t} = \sum_{k=1}^{m} \frac{Sales_{i,j,t}}{Sales_{i,t}} \times \frac{Robots_{j,t}^{EU}}{Workers_{j,t}^{EU}},$$ where $Robots_{j,t}^{EU}$ and $Workers_{j,t}^{EU}$ represent the counts of industrial robots and workers, respectively, in industry j during year t across EU countries including Denmark, France, Finland, Italy, and Sweden.
Innovation Variable	
Tech Similarity	We set a rolling time window of 5 years and calculate the ratio of patents in each Cooperative Patent Classification (CPC) for each company to obtain a series of vectors $T_i = (T_{i1},\ T_{i2},\ \ldots,\ T_{in})$,where T_{ij} represents the ratio of patents in patent classification j to all patents invented by firm i. Then we calculate Technological Shift as follows: $$Technological\ Shift_{i,t} = \frac{(T_{i,t}T'_{i,t-5})}{(T_{i,t}T'_{i,t})^{\frac{1}{2}}(T_{i,t-5}T'_{i,t-5})^{\frac{1}{2}}}$$ where $T_{i,t}$ and $T_{i,t-5}$ is firm i's vectors of patent distribution in year t and t-5, respectively.

Variable	Definition
BERT Similarity	We employ a pre-trained BERT model, PatentBERT, to encode each patent document into a 768-dimensional continuous vector, capturing a dense semantic representation of its content. For each firm in each year, we then calculate a 768-dimensional continuous vector, which represents the arithmetic mean of vectors for all patents filed by the firm that year. Subsequently, we apply the same methodology used in calculating Tech Similarity to obtain the cosine similarity of the firm's patent distributions between year t and $t+5$.
R&D/Total Asset	R&D expenditure in a given year divided by lagged total assets.
R&D/Sales	R&D expenditure in a given year divided by lagged total sales.
Patent Number	The total number of patent applications filed (and eventually granted by the CNIPA) by a firm in a given year.
Citation-Deflated Patent Number	We calculate the weighted average of each patent, where the weight is calculated as follows: $$Citation\text{-}Deflated\ Weight_p = \sum_{j=1}^{m} \frac{Total\ Citation_i}{m} \times \frac{1}{Average\ CPC\ Citation_j}$$ Where, $Total\ Citation_i$ denotes the total citations for patent p over the next ten years, m denotes the number of CPC fields to which a patent belongs, $Average\ CPC\ Citation_j$ denotes CPC field j denotes the average number of citations received by CPC field j from all n CPC fields over the next ten years.
Data Generation	We calculate the frequency of data-generation-related keywords in each company's annual 10-K report and normalize it by the total word count of the report for each year.

Variable	Definition
Firm Characteristics Variable	
Ln (Assets)	Logged value of the book value of total assets.
Leverage	Sum of long- and short-term debt (dltt+dlc) divided by the book value of total assets(at).
Ln(Age)	Logged value of one plus the difference between the year under investigation and the first year the firm appears on the CRSP tapes.
ROA	The ratio of the firm's operating income before depreciation (oibdp) divided by the lagged book value of total assets (at).
Tangibility	The ratio of the firm's tangible Assets(1-intan) divided by the book value of total assets(at).
Panel B: Patent Level	
AI Distance	We calculate the AI distance based on the citation network. According to the minimum number of direct and indirect citations of AI patents, when a patent directly cites AI patents, the patent distance is 1; when a patent cites AI patents by citing other patents, the patent distance n is obtained by adding up the times that patents indirectly citing. The AI patent's patent distance is set to zero.

Variable	Definition
AI Closeness	Each year t, we construct a citation distribution vector $C_{AI,t}$ for a representative AI patent across USPTO technology classes, normalizing citation counts by class to capture its technological engagement. For any patent p filed in year t, we similarly define a citation distribution vector $C_{p,t}$. The "AI Closeness" of patent p is then measured as the cosine similarity between $C_{AI,t}$ and $C_{p,t}$: $$AI\ Closeness_{p,t} = \frac{C\prime_{AI,t} C_{p,t}}{\|C_{AI,t}\|_2 \|C_{p,t}\|_2}$$ where $\|x\|_2$ represents the L_2 norm, ensuring vector normalization.
Originality	We calculate Originality as follows: $$Originality_p = 1 - \sum_{j=1}^{N} q_{pj}^2$$ where q_{pj} is the share of patents that patent p cites in technology class j.
Generality	We calculate Generality as follows: $$Generality_p = 1 - \sum_{j=1}^{N} s_{pj}^2$$ where s_{pj} is share of citations that patent p receives from technology class j within 5 years after patent p is issued.
Team Size	We measure a patent's team size as the number of inventors.
Labor Input	We calculate each patent's labor input that deflates the team size's double counting.

Variable	Definition
Inventor Originality	We measure an inventor's originality as follows. $$Originality_{i,t} = 1 - \sum_{j=1}^{N} s_{ji,t}^2$$ where $s_{ji,t}$ is the fraction of patents by inventor i that falls into technology j in the past five years.

Table A2: Summary Statistics

This table presents summary statistics for key variables used in our analysis. Panel A includes firm-level variables, encompassing robot-related, innovation-related, and firm-characteristic-related metrics, while Panel B provides a summary of patent-level variables. Definitions for all variables are available in the Appendix (see Table A1).

Variable	N	Mean	S.D.	Min	Max
Panel A: Firm-level					
Robot-related					
Robot Exposure	20,962	0.000	1.000	-0.529	6.410
Robot Exposure (IV)	20,962	13.239	20.356	0.000	139.106
Innovation-related					
Tech Similarity	20,962	0.906	0.157	0.000	1.000
BERT Similarity	18,003	0.760	0.349	0.000	1.000
Citation-Deflated BERT Similarity	18,003	0.758	0.349	0.000	1.000
R&D/Total Assets	20,962	0.462	0.311	0.000	1.112
R&D/Sales	20,962	0.482	0.629	0.000	7.070
Patent Number	20,962	30.193	214.830	0.000	8,878.000
Citation-Deflated Patent Number	20,962	2.420	13.712	0.000	495.451
Data Generation	20,962	0.000	1.000	-0.761	3.643
Firm-characteristics-related					
Ln(Assets)	20,962	2.954	0.731	0.000	4.111
Leverage	20,962	6.200	2.114	0.619	11.790
Ln(Age)	20,962	0.943	0.633	0.000	4.464
ROA	20,962	0.480	0.336	0.017	15.481
Tangibility	20,962	0.825	0.185	0.235	1.000
Panel B: Patent-level					
AI Distance	661,240	5.561	6.753	0.000	24.000
AI Closeness	661,240	0.271	0.230	0.000	1.000
Originality	608,205	0.369	0.271	0.000	0.945
Generality	365,241	0.128	0.204	0.000	0.921
Team Size	651,475	2.992	1.960	1.000	60.000
Labor Input	651,475	1.530	1.448	0.003	40.667
Inventor Originality	621,621	0.335	0.258	0.000	0.929

Table A3: SIC15 Categories and Corresponding SIC4 Industry Classifications

This table presents the classification of industries based on the 4-digit Standard Industrial Classification codes, aggregated into 15 broad categories (SIC15). Each category includes specific ranges of SIC codes that represent different types of industries.

SIC15	SIC15 Industry Description	SIC4
1	Agriculture, hunting, forestry, and fishing	0100–0999
2	Mining and quarrying	1000–1499
3	Food products and beverages; tobacco products	2000–2199
4	Textiles, leather, wearing apparel	2200–2399, 3100–3199
5	Wood and wood products, paper and paper products, publishing and printing	2400–2599, 2600–2799
6	Plastic and chemical products, glass, ceramics, stone, mineral products	2800–3099, 3200–3299
7	Basic metals and metal products	3300–3399, 3400–3499
8	Metal machinery	3500–3599
9	Electrical/electronics	3600–3699, 3800–3899
10	Automotive and other transport equipment	3700–3799
11	All other manufacturing branches	3900–3999
12	Electricity, gas and water supply	4900–4999
13	Construction	1500–1799
14	Education, research and development	8200–8299, 8700–8799
15	All other non-manufacturing branches	4000–4899, 5000–6999, 7000–8199, 8300–8699, 8800–9999

Table A4: SIC15 Categories and Corresponding IFR Industry Classifications

This table provides the mapping between SIC15 categories and their corresponding IFR industry classifications, along with descriptions of each category.

SIC15	IFR Industry	IFR Industry Description
1	A-B	Agriculture, forestry, fishing
2	C	Mining and quarrying
3	D10-D12	Food and beverages
4	D13-D15	Textiles
5	D16-D18	Wood and furniture; paper
6	D19-D23	Plastic and chemical products; glass, ceramics, stone, mineral products
7	D24-D25	Basic metals; metal products
8	D28	Industrial machinery
9	D26-D27	Electronic components/devices; electrical machinery
10	D29-D30	Automotive; other vehicles
11	91	All other manufacturing branches
12	E	Electricity, gas, water supply
13	F	Construction
14	P	Education/research/development
15	90	All other non-manufacturing branches

Table A5: SIC15 Categories and Corresponding EU KLEMS Industry Classifications

This table provides the mapping between SIC15 categories and their corresponding EU KLEMS industry classifications, along with descriptions of each category.

SIC15	EU KLEMS Industry	EU KLEMS Industry Description
1	A	Agriculture, forestry and fishing
2	B	Mining and quarrying
3	C10-C12	Manufacture of food products; beverages and tobacco products
4	C13-C15	Manufacture of textiles, wearing apparel, leather, etc.
5	C16-C18	Manufacture of wood, paper, printing and reproduction
6	C19-C23	Manufacture of coke and refined petroleum products; chemicals, pharmaceutical and plastic products
7	C24-C25	Manufacture of basic metals and fabricated metal products
8	C28	Manufacture of machinery and equipment
9	C26-C27	Manufacture of electrical equipment
10	C29-C30	Manufacture of motor vehicles, trailers, semi-trailers and of other transport equipment
11	C31-C33	Manufacture of furniture; jewellery, toys; repair and installation
12	D-E	Electricity, gas, steam; water supply, sewerage, waste management
13	F	Construction
14	M-Q	Professional, scientific and technical activities; education, human health and social work activities
15	G-L, R-U	All other non-manufacturing branches such as wholesale and retail trade; transportation and storage, etc.

Table A6: 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 determine the frequency of patent abstracts for both groups relative to the baseline year of 2004. The relative frequencies are then calculated (after log-transformation) to highlight significant differences between the two groups.

Year	2010		2015		2019	
ID	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 A7: Keyword List for Defining Data Generation

This table presents the keyword list used to define firm-level data generation. We use the frequency of these data-generation-related keywords in 10-K filings as a proxy for data generation. Column 1 provides references to a broad range of literature, while the left columns display the keyword list. In counting keyword frequencies, distinctions in case sensitivity and singular/plural forms are disregarded.

References	Keywords
Brynjolfsson and Kim (2011)	Data asset
Chen et al. (2012)	Rich information
McAfee and Brynjolfsson (2012)	Digital technology
Tambe (2014)	Data mart
Brynjolfsson and McElheran (2016)	Data sharing
Brynjolfsson and Mitchell (2017)	Information provision
Begenau et al. (2018)	Intensive information
Agrawal et al. (2019)	BI&A (business intelligence & analytics)
Bajari et al. (2019)	Data gathering
Chen et al. (2019)	Data storage
Farboodi et al. (2019)	Database management
Fuster et al. (2019)	Data driven
Goldstein et al. (2019)	Rich data
Zhu (2019)	Business data
Wu et al. (2020)	Data mining
	Data generation
	Data accumulation
	Intangible asset
	Information based
	Electronic data

Table A8: Comparison of AI patents to non-AI patents

This table presents the results of the OLS regression, analyzing the relative costs associated with AI patents in comparison to non-AI patents. Specifically, the following model is estimated:

$$y_{i,p,t} = \alpha_{i,t} + \beta \times Dummy(AI\ patent)_{i,p,t} + \varepsilon_{i,p,t},$$

where $y_{i,p,t}$ represents the cost of firm i's patent p 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. Columns (1)–(8) report the estimated coefficients of $Dummy(AI\ patent)$ for *Team Size*, *Labor Input*, *Inventor Originality*, and *Patent Value* respectively, in year t. We test two types of fixed effects: one by including firm and year fixed effects in the regression, and the other by including firm-year fixed effects. 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 is at the patent-firm-year level. Robust standard errors are provided 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	651,475	651,475	651,475	651,475	621,621	621,621	659,189	657,128

Table A9: The impact of Robot Exposure on Innovation Shifts Over the Subsequent Years

This table presents the second-stage results of the 2SLS estimation, examining the impact of robot exposure on innovation shifts over the subsequent years. Panel A displays the effects on firms' innovation similarity at firm-year level, including *Tech Similarity* and *BERT Similarity*. Panel B presents the impact on AI-directed innovation at patent-firm-year level, including *AI Distance* and *AI Closeness*. Panel C presents the impact on innovation inputs at firm-year level, including *R&D/Total Assets* and *R&D/Sales*. *Tech Similarity* is measured by the cosine similarity of a firm's patent distribution over a five-year period. *BERT Similarity* is calculated as the cosine similarity of the firm's 768-dimensional continuous vector, generated by the PatentBERT model, over five years. *AI Distance* represents the minimum number of steps required for a firm's patent to reach AI patents within the citation network in the given year. *AI Closeness* is calculated as the cosine similarity between the citation distribution vectors of a firm's patent and AI patents in the given year. *Robot Exposure* quantifies the degree of robot exposure experienced by the firm in the given year, standardized to have a mean of zero and a standard deviation of one. Columns (1)–(6) report the estimated coefficients for the effects over the subsequent one to six years, 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. Robust standard errors are provided in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Model	IV					
Period	t+1	t+2	t+3	t+4	t+5	t+6
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: Innovation Similarity					
Level	Firm-Year					
Variable	Tech Similarity					
Robot Exposure	-0.213*	-0.263*	-0.169	-0.035	0.035	-0.011
	(0.123)	(0.158)	(0.137)	(0.131)	(0.128)	(0.136)
Observations	18,088	15,908	13,945	12,198	10,632	9,252
Variable	BERT Similarity					
Robot Exposure	-0.722**	-0.812*	-0.420	-0.071	0.083	0.029
	(0.363)	(0.466)	(0.259)	(0.177)	(0.154)	(0.170)
Observations	15,625	13,793	12,118	10,607	9,259	8,050
	Panel B: AI-Directed Innovation					
Level	Patent-Firm-Year					
Variable	AI Distance					
Robot Exposure	-0.581***	-0.556***	-0.427***	-0.287	-0.014	-0.160
	(0.124)	(0.140)	(0.155)	(0.176)	(0.188)	(0.196)
Observations	607,772	561,545	517,043	472,706	428,740	391,768
Variable	AI Closeness					
Robot Exposure	0.084***	0.087***	0.091***	0.102***	0.086***	0.076***
	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)	(0.005)
Observations	607,772	561,545	517,043	472,706	428,740	391,768

(Continued)

Model	IV					
Period	t+1	t+2	t+3	t+4	t+5	t+6
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel C: Innovation Input					
Level	Firm-Year					
Variable	R&D/Total Assets					
Robot Exposure	0.207**	0.182*	0.088	0.057	0.056	0.063
	(0.102)	(0.109)	(0.097)	(0.097)	(0.073)	(0.079)
Observations	18,088	15,908	13,945	12,198	10,632	9,252
Variable	R&D/Sales					
Robot Exposure	0.684**	0.784**	0.672**	0.576*	0.307	0.109
	(0.296)	(0.382)	(0.330)	(0.336)	(0.219)	(0.188)
Observations	18,088	15,908	13,945	12,198	10,632	9,252
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: Robustness Checks of Controlling for Industry-level Price Changes

This table presents the results of robustness checks by controlling for industry-level Producer Price Index. Panels A and B display the second-stage and first-stage robustness results of the 2SLS estimation, respectively. In Panel A, the dependent variables in Columns (1)–(6) include *Tech Similarity*, *BERT Similarity*, *AI Distance*, *AI Closeness*, *R&D/Total Assets*, and *R&D/Sales*. In Panel B, the dependent variable is *Robot Exposure*, which quantifies the degree of robot exposure faced by the firm in a given year, standardized to have a mean of zero and a standard deviation of one. Additionally, *Robot Exposure IV* is measured using the ratio of the aging population to the workforce in European countries. 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 for testing the effect on *AI Distance* and *AI Closeness* are at the patent-firm-year level, while the remaining variables are at the firm-year level. All regressions incorporate firm-fixed and year-fixed effects. Robust standard errors are provided 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)
	Controlling for Industry-level PPI					
	Panel A: The Second Stage of 2SLS Estimates					
Variable	Tech Similarity	BERT Similarity	AI Distance	AI Closeness	R&D/Total Assets	R&D/Sales
Robot Exposure	-0.179*	-0.488**	-1.039***	0.075***	0.199**	0.556***
	(0.092)	(0.241)	(0.130)	(0.004)	(0.082)	(0.215)
PPI	-0.000	0.000	-0.023***	0.002***	-0.000***	0.001***
	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)	(0.000)
	Panel B: The First Stage of 2SLS Estimates					
Variable	Robot Exposure					
Robot Exposure IV	0.003***	0.003***	0.014***	0.014***	0.003***	0.003***
	(0.001)	(0.001)	(0.000)	(0.000)	(0.001)	(0.001)
PPI	-0.000**	-0.000**	-0.001***	-0.001***	-0.000**	-0.000**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Observations	20,251	17,537	644,199	644,199	20,251	20,251
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: Robustness Checks of Controlling for Industry-Year Fixed Effects

This table presents the results of robustness checks by controlling for industry-year fixed effects, where industry is defined at the SIC 2-digit level. Panels A and B display the second-stage and first-stage robustness results of the 2SLS estimation, respectively. In Panel A, the dependent variables in Columns (1)–(6) include *Tech Similarity*, *BERT Similarity*, *AI Distance*, *AI Closeness*, *R&D/Total Assets*, and *R&D/Sales*. In Panel B, the dependent variable is *Robot Exposure*, which quantifies the degree of robot exposure faced by the firm in a given year, standardized to have a mean of zero and a standard deviation of one. Additionally, *Robot Exposure IV* is measured using the ratio of the aging population to the workforce in European countries. 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 for testing the effect on *AI Distance* and *AI Closeness* are at the patent-firm-year level, while the remaining variables are at the firm-year level. All regressions incorporate firm-fixed and industry-year-fixed effects. Robust standard errors are provided 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)
	Controlling for Industry-Year Fixed Effects					
	Panel A: The Second Stage of 2SLS Estimates					
Variable	Tech Similarity	BERT Similarity	AI Distance	AI Closeness	R&D/Total Assets	R&D/Sales
Robot Exposure	-0.194*	-0.489*	-4.527***	0.091***	0.195*	0.493*
	(0.117)	(0.273)	(1.033)	(0.023)	(0.105)	(0.261)
	Panel B: The First Stage of 2SLS Estimates					
Variable	Robot Exposure					
Robot Exposure IV	0.004***	0.004***	0.004***	0.004***	0.004***	0.004***
	(0.001)	(0.001)	(0.000)	(0.000)	(0.001)	(0.001)
Observations	20,699	17,798	660,963	660,963	20,699	20,699
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
SIC2×Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table A12: Robustness Checks of pre-trend Test

This table presents the results of robustness checks, confirming that there are no pre-existing trends in prior innovation activities. We assess the effect of robot exposure during the 2004–2019 period on firms' innovation shifts and innovation productivity between 1988 and 2003. Panels A, B, and C display the second-stage robustness results of the 2SLS estimation, with Panel A focusing on firms' innovation similarity, including *Tech Similarity* and *BERT Similarity*, Panel B focusing on AI-directed innovation, including *AI Distance* and *AI Closeness*, and Panel C examining the effects on firms' innovation inputs, including $R\&D/Total\ Assets$ and $R\&D/Sales$. *Robot Exposure* quantifies the degree of robot exposure faced by the firm in a given year, standardized to have a mean of zero and a standard deviation of one. Columns (1)–(3) report the coefficients of robot exposure on the first independent variable for years t, $t+1$, and $t+2$, respectively. Columns (4)–(6) present the coefficients for the second independent variable for the same 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. The regression data for all panels are structured at the firm-year level. In Panel B, we employ the citation-weighted measure to calculate firm-level AI distance and closeness, ensuring alignment with the variables from prior periods. All regressions incorporate firm-fixed and year-fixed effects. Robust standard errors are provided in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Model	IV					
Period	t	t+1	t+2	t	t+1	t+2
	(1)	(2)	(3)	(4)	(5)	(6)
	Pre-trend Tests					
	Panel A: Innovation Similarity (1988-2003)					
Variable	Tech Similarity			BERT Similarity		
Robot Exposure (2004-2019)	0.086	0.134	0.057	-0.533	1.034	9.617
	(0.089)	(0.119)	(0.157)	(0.455)	(2.593)	(99.790)
Observations	17,392	14,943	13,062	3,772	3,623	3,522
	Panel B: AI-Directed Innovation (1988-2003)					
Variable	AI Distance			AI Closeness		
Robot Exposure (2004-2019)	-0.929	-1.787	-5.547	-0.009***	-0.008**	-0.006**
	(2.417)	(3.160)	(4.556)	(0.003)	(0.003)	(0.003)
Observations	20,724	18,014	15,846	20,724	18,014	15,846
	Panel C: Innovation Input (1988-2003)					
Variable	R&D/Total Assets			R&D/Sales		
Robot Exposure (2004-2019)	0.560	0.581	1.017	1.148	1.635	2.461
	(0.541)	(0.684)	(1.673)	(1.043)	(1.801)	(4.007)
Observations	11,017	9,500	8,346	11,017	9,500	8,346
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table A13: Robustness Checks of Controlling for the Change in the Dependent Variables

This table presents the results of robustness checks by directly controlling for changes in firms' innovation activities between 1988 and 2003. Panels A and B display the second-stage and first-stage robustness results of the 2SLS estimation, respectively. In Panel A, the dependent variables in Columns (1)–(6) include *Tech Similarity*, *BERT Similarity*, *AI Distance*, *AI Closeness*, *R&D/Total Assets*, and *R&D/Sales*. In Panel B, the dependent variable is *Robot Exposure*, which quantifies the degree of robot exposure faced by the firm in a given year, standardized to have a mean of zero and a standard deviation of one. Additionally, *Robot Exposure IV* is measured using the ratio of the aging population to the workforce in European countries. 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 for all panels are structured at the firm-year level. We employ the citation-weighted measure to calculate firm-level AI distance and closeness, ensuring alignment with the variables from prior periods. All regressions incorporate firm-fixed and year-fixed effects. Robust standard errors are provided 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)
	Controlling for the Change in the Dependent Variable (1988–2003)					
	Panel A: The Second Stage of 2SLS Estimates					
Variable	Tech Similarity	BERT Similarity	AI Distance	AI Closeness	R&D/Total Assets	R&D/Sales
Robot Exposure	-0.186**	-0.460**	-5.793*	0.004**	0.192**	0.633***
	(0.093)	(0.229)	(3.265)	(0.002)	(0.080)	(0.232)
Change in Dep. Var. (1988-2003)	-0.009	0.002	-0.048***	0.031	-0.000	-0.018*
	(0.011)	(0.023)	(0.016)	(0.057)	(0.011)	(0.010)
	Panel B: The First Stage of 2SLS Estimates					
Variable	Robot Exposure					
Robot Exposure IV	0.003***	0.003***	0.003***	0.003***	0.003***	0.003***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Change in Dep. Var. (1988-2003)	0.013	0.011	-0.002**	-0.894	-0.064**	0.002
	(0.028)	(0.030)	(0.001)	(3.216)	(0.026)	(0.008)
F Statistics	23.95	18.40	22.91	23.71	24.15	23.96
Observations	20,815	17,915	20,552	20,552	20,815	20,815
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table A14: Robustness Checks of Alternative Instrument Variables: 6 European Countries

This table presents the results of robustness checks, using alternative instrument variables measures to examine the impact of robot exposure on firms' innovation. We replace the industry-level instrument with the average automation index of the five European countries plus Germany. Panels A, B, and C display the second-stage robustness results of the 2SLS estimation, with Panel A focusing on firms' innovation similarity, including *Tech Similarity* and *BERT Similarity*, Panel B focusing on AI-directed innovation, including *AI Distance* and *AI Closeness*, and Panel C examining the effects on firms' innovation inputs, including *R&D/Total Assets* and *R&D/Sales*. *Robot Exposure* quantifies the degree of robot exposure faced by the firm in a given year, standardized to have a mean of zero and a standard deviation of one. Columns (1)–(3) report the coefficients of robot exposure on the first independent variable for years t, $t+1$, and $t+2$, respectively. Columns (4)–(6) present the coefficients for the second independent variable for the same 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. The regression data for Panels A and C are at the firm-year level, while Panel B is at the patent-firm-year level. All regressions incorporate firm-fixed and year-fixed effects. Robust standard errors are provided in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Model	IV					
Period	t	t+1	t+2	t	t+1	t+2
	(1)	(2)	(3)	(4)	(5)	(6)
	6 European Countries IV					
	Panel A: Innovation Similarity					
Variable	Tech Similarity			BERT Similarity		
Robot Exposure	-0.108*	-0.116*	-0.108	-0.368**	-0.444**	-0.340*
	(0.059)	(0.069)	(0.076)	(0.163)	(0.198)	(0.181)
Observations	18,090	15,633	13,676	15,627	13,552	11,879
	Panel B: AI-Directed Innovation					
Variable	AI Distance			AI Closeness		
Robot Exposure	-0.803***	-0.389**	-0.234	0.075***	0.087***	0.081***
	(0.140)	(0.159)	(0.194)	(0.004)	(0.004)	(0.005)
Observations	607,772	554,683	510,443	607,772	554,683	510,443
	Panel C: Innovation Input					
Variable	R&D/Total Assets			R&D/Sales		
Robot Exposure	0.120**	0.111**	0.092	0.389***	0.357**	0.315**
	(0.051)	(0.055)	(0.059)	(0.147)	(0.157)	(0.153)
Observations	18,090	15,633	13,676	18,090	15,633	13,676
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table A15: Robustness Checks of Alternative Instrument Variables: 9 European Countries

This table presents the results of robustness checks, using alternative instrument variables measures to examine the impact of robot exposure on firms' innovation. We replaces the industry-level instrument with the average automation index of all nine advanced automation countries in Europe. These nine countries are Denmark, Finland, France, Germany, Italy, Norway, Spain, Sweden, and the United Kingdom. Panels A, B, and C display the second-stage robustness results of the 2SLS estimation, with Panel A focusing on firms' innovation similarity, including *Tech Similarity* and *BERT Similarity*, Panel B focusing on AI-directed innovation, including *AI Distance* and *AI Closeness*, and Panel C examining the effects on firms' innovation inputs, including *R&D/Total Assets* and *R&D/Sales*. *Robot Exposure* quantifies the degree of robot exposure faced by the firm in a given year, standardized to have a mean of zero and a standard deviation of one. Columns (1)–(3) report the coefficients of robot exposure on the first independent variable for years t, $t+1$, and $t+2$, respectively. Columns (4)–(6) present the coefficients for the second independent variable for the same 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. The regression data for Panels A and C are at the firm-year level, while Panel B is at the patent-firm-year level. All regressions incorporate firm-fixed and year-fixed effects. Robust standard errors are provided in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Model	IV					
Period	t	t+1	t+2	t	t+1	t+2
	(1)	(2)	(3)	(4)	(5)	(6)
	9 European Countries IV					
	Panel A: Innovation Similarity					
Variable	Tech Similarity			BERT Similarity		
Robot Exposure	-0.210**	-0.256*	-0.324	-0.499*	-0.782*	-0.887
	(0.106)	(0.148)	(0.197)	(0.265)	(0.433)	(0.546)
Observations	20,815	18,090	15,909	17,915	15,627	13,794
	Panel B: AI-Directed Innovation					
Variable	AI Distance			AI Closeness		
Robot Exposure	-1.022***	-0.634***	-0.570***	0.083***	0.087***	0.087***
	(0.118)	(0.120)	(0.138)	(0.003)	(0.003)	(0.004)
Observations	661,018	607,468	561,097	661,018	607,468	561,097
	Panel C: Innovation Input					
Variable	R&D/Total Assets			R&D/Sales		
Robot Exposure	0.166**	0.190*	0.169	0.638**	0.712**	0.818*
	(0.084)	(0.109)	(0.118)	(0.252)	(0.335)	(0.436)
Observations	20,815	18,090	15,909	20,815	18,090	15,909
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table A16: Robustness Checks of Alternative Instrument Variables: Filter Time Fixed Effect

This table presents the results of robustness checks, using alternative instrument variables measures to examine the impact of robot exposure on firms' innovation. We regress robot exposure on time fixed effects and denote the residuals from this regression as the alternative measure. Panels A, B, and C display the second-stage robustness results of the 2SLS estimation, with Panel A focusing on firms' innovation similarity, including *Tech Similarity* and *BERT Similarity*, Panel B focusing on AI-directed innovation, including *AI Distance* and *AI Closeness*, and Panel C examining the effects on firms' innovation inputs, including $R\&D/Total\ Assets$ and $R\&D/Sales$. *Robot Exposure* quantifies the degree of robot exposure faced by the firm in a given year, standardized to have a mean of zero and a standard deviation of one. Columns (1)–(3) report the coefficients of robot exposure on the first independent variable for years t, $t+1$, and $t+2$, respectively. Columns (4)–(6) present the coefficients for the second independent variable for the same 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. The regression data for Panels A and C are at the firm-year level, while Panel B is at the patent-firm-year level. All regressions incorporate firm-fixed and year-fixed effects. Robust standard errors are provided in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Period	t	t+1	t+2	t	t+1	t+2
	(1)	(2)	(3)	(4)	(5)	(6)
	IV with Time Fixed Effects Filtered					
	Panel A: Innovation Similarity					
Model	IV					
Variable	Tech Similarity			BERT Similarity		
Robot Exposure	-0.187**	-0.212*	-0.260*	-0.460**	-0.716**	-0.803*
	(0.093)	(0.122)	(0.156)	(0.229)	(0.358)	(0.458)
Observations	20,815	18,090	15,909	17,915	15,627	13,794
	Panel B: AI-Directed Innovation					
Model	IV					
Variable	AI Distance			AI Closeness		
Robot Exposure	-1.004***	-0.581***	-0.556***	0.079***	0.084***	0.087***
	(0.124)	(0.124)	(0.140)	(0.004)	(0.004)	(0.004)
Observations	661,018	607,772	561,545	661,018	607,772	561,545
	Panel C: Innovation Input					
Model	IV					
Variable	R&D/Total Assets			R&D/Sales		
Robot Exposure	0.192**	0.206**	0.181*	0.626***	0.680**	0.780**
	(0.080)	(0.101)	(0.108)	(0.230)	(0.293)	(0.378)
Observations	20,815	18,090	15,909	20,815	18,090	15,909
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table A17: Falsification Tests of Shift-share IV

This table presents the falsification tests results of the shift-share IV following Borusyak et al. [2022]. We conduct placebo tests using industry-level and firm-level characteristic variables to verify the randomness of the shock from Europe. And in order to obtain valid exposure-robust standard errors, we estimate these coefficients with equivalent industry-level regressions. Panels A, B, and C display these coefficients. Panel A focusing on industry-level potential confounders. Following Acemoglu and Restrepo [2020], we select industry-level control variables, including the log of real wages (Ln(Wage)), wage per worker (Wage/Worker), the log of employment (Ln(Worker)), the share of high-tech equipment in total investment (High-tech/Invest), the ratio of input to value-added (Input/Value-added), and the ratio of workers to value-added (Worker/Value-added). Panel B examining the potential effects of firm-level variables, including the basic control variables and a dummy variable indicating whether the firm uses an ES system. And Panel C focusing on pre-trend variables from 1988 to 2003. *Robot Exposure* quantifies the degree of robot exposure faced by the firm in a given year, standardized to have a mean of zero and a standard deviation of one. The sample period spans 2004–2019. The regression data for Panels A, B, and C are at the industry-year level. All regressions incorporate year-fixed effects. Exposure-robust standard errors are provided in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Model	IV					
Period	t+1					
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: Industry-level Balance Test					
Variable	Ln(Wage)	Wage/Worker	Ln(Worker)	High-tech/Invest	Input/Value-added	Worker/Value-added
Robot Exposure	-0.359	2.510	-0.350	-0.321	3.311	0.067
	(0.289)	(1.675)	(0.286)	(0.233)	(2.431)	(0.047)
No. of industry-periods	195	195	195	195	195	195
	Panel B: Firm-level Balance Test					
Variable	ES	Ln(Assets)	ROA	Leverage	Tangibility	Ln(Age)
Robot Exposure	0.444	3.023	-1.120*	0.474	-0.137	1.296***
	(0.271)	(1.843)	(0.551)	(0.370)	(0.124)	(0.416)
No. of industry-periods	195	195	195	195	195	195
	Panel C: Firm-level pre-trends Test (1988-2003)					
Variable	Tech Similarity	BERT Similarity	AI Distance	AI Closeness	R&D/Total Assets	R&D/Sales
Robot Exposure (2004-2019)	-0.078	1.491	-0.876	-0.003	0.000	0.203
	(0.181)	(3.208)	(1.660)	(0.002)	(0.022)	(0.456)
No. of industry-periods	48	48	188	188	48	48
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table A18: Robustness Checks of Alternative Instrument Variables: Shift-share IV

This table presents the results of robustness checks, employing the shift-share instrumental variable to examine the impact of robot exposure on firms' innovation, following the research design proposed by Borusyak et al. [2022]. We replace the time-varying variable in the original instrumental variable calculation, which represents the annual proportion of robot usage to employment across EU countries, with the fixed average proportion of robot usage to employment calculated from the years 2000 to 2004. We estimate these coefficients with equivalent industry-level regressions in order to obtain valid exposure-robust standard errors. Panels A, B, and C display these coefficients, with Panel A focusing on firms' innovation similarity, including *Tech Similarity* and *BERT Similarity*, Panel B focusing on AI-directed innovation, including *AI Distance* and *AI Closeness*, and Panel C examining the effects on firms' innovation inputs, including *R&D/Total Assets* and *R&D/Sales*. *Robot Exposure* quantifies the degree of robot exposure faced by the firm in a given year, standardized to have a mean of zero and a standard deviation of one. Columns (1)–(3) report the coefficients of robot exposure on the first independent variable for years t, $t+1$, and $t+2$, respectively. Columns (4)–(6) present the coefficients for the second independent variable for the same years. The sample period spans 2004–2019. The regression data for Panels A, B, and C are at the industry-year level. All regressions incorporate year-fixed effects. Exposure-robust standard errors are provided in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Model	IV					
Period	t	t+1	t+2	t	t+1	t+2
	(1)	(2)	(3)	(4)	(5)	(6)
	Shift-share IV					
	Panel A: Innovation Similarity					
Variable	Tech Similarity			BERT Similarity		
Robot Exposure	-0.505	-0.544	-0.528	-1.648	-1.710	-1.756
	(0.909)	(0.995)	(0.973)	(2.685)	(2.855)	(2.927)
No. of industry-periods	208	195	182	208	195	182
	Panel B: AI-Directed Innovation					
Variable	AI Distance			AI Closeness		
Robot Exposure	-1.582**	-1.546**	-1.559**	0.050**	0.051**	0.052**
	(5.677)	(5.565)	(5.469)	(0.021)	(0.021)	(0.020)
No. of industry-periods	208	195	182	208	195	182
	Panel C: Innovation Input					
Variable	R&D/Total Assets			R&D/Sales		
Robot Exposure	0.932*	0.917*	0.913*	1.545**	1.507**	1.486**
	(0.496)	(0.462)	(0.448)	(0.559)	(0.512)	(0.489)
No. of industry-periods	208	195	182	208	195	182
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table A19: Robustness Checks of Alternative Robot Exposure Measures: LASSO

This table presents the results of robustness checks, using alternative robot exposure measures to examine the impact of robot exposure on firms' innovation. We apply a two-stage regression process to filter and re-weight keywords, following Cong et al. [2019], and then standardize the measure to have a mean of zero and a standard deviation of one, denoted as *Robot Exposure A*. Panels A, B, and C display the second-stage robustness results of the 2SLS estimation, with Panel A focusing on firms' innovation similarity, including *Tech Similarity* and *BERT Similarity*, Panel B focusing on AI-directed innovation at the patent-firm-year level, including *AI Distance* and *AI Closeness*, and Panel C examining the effects on firms' innovation inputs, including *R&D/Total Assets* and *R&D/Sales*. Panel D presents the Poisson regression robustness results for innovation outputs, including *Patent Number* and *Citation-Deflated Patent Number*. Columns (1)–(3) report the coefficients of robot exposure on the first independent variable for years t, $t+1$, and $t+2$, respectively. Columns (4)–(6) present the coefficients for the second independent variable for the same 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. The regression data for Panels A, C, and D are at the firm-year level, while Panel B is at the patent-firm-year level. All regressions incorporate firm-fixed and year-fixed effects. Robust standard errors are provided in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Period	t	t+1	t+2	t	t+1	t+2
	(1)	(2)	(3)	(4)	(5)	(6)
	LASSO-Method Robot Exposure					
	Panel A: Innovation Similarity					
Model	IV					
Variable	Tech Similarity			BERT Similarity		
Robot Exposure A	-0.760	-0.828	-1.013	-1.371*	-1.979	-2.346
	(0.482)	(0.611)	(0.795)	(0.807)	(1.209)	(1.684)
Observations	20,815	18,090	15,909	17,915	15,627	13,794
	Panel B: AI-Directed Innovation					
Model	IV					
Variable	AI Distance			AI Closeness		
Robot Exposure A	-8.071***	-6.051***	-23.719**	0.633***	0.871***	3.705***
	(1.059)	(1.340)	(9.248)	(0.041)	(0.065)	(1.111)
Observations	661,018	607,772	561,545	661,018	607,772	561,545
	Panel C: Innovation Input					
Model	IV					
Variable	R&D/Total Assets			R&D/Sales		
Robot Exposure A	0.782*	0.805	0.704	2.545*	2.659	3.038
	(0.446)	(0.546)	(0.569)	(1.361)	(1.688)	(2.186)
Observations	20,815	18,090	15,909	20,815	18,090	15,909
	Panel D: Innovation Output					
Model	Possion			IV		
Variable	Citation-Deflated Patent Number			Citation-Deflated Patent Value		
Robot Exposure A	-0.052*	-0.038	-0.054***	-4.053	-4.549	-6.200

Table A20: Robustness Checks of Alternative Robot Exposure Measures: IFR

This table presents the results of robustness checks, using alternative robot exposure measures to examine the impact of robot exposure on firms' innovation. We use the ratio of robots to employment and then standardize the measure to have a mean of zero and a standard deviation of one, denoted as *Robot Exposure B*. Specifically, we first measure robot usage using data from the IFR. Second, we collect employment data from the EU KLEMS database. To construct firm-level robot exposure, we combine our industry-level robot exposure dataset with the Compustat-Segment dataset. Panels A, B, and C display the second-stage robustness results of the 2SLS estimation, with Panel A focusing on firms' innovation similarity, including *Tech Similarity* and *BERT Similarity*, Panel B focusing on AI-directed innovation, including *AI Distance* and *AI Closeness*, and Panel C examining the effects on firms' innovation inputs, including $R\&D/Total\ Assets$ and $R\&D/Sales$. Panel D presents the Poisson regression robustness results for innovation outputs, including *Patent Number* and *Citation-Deflated Patent Number*. Columns (1)–(3) report the coefficients of robot exposure on the first independent variable for years t, $t+1$, and $t+2$, respectively. Columns (4)–(6) present the coefficients for the second independent variable for the same 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. The regression data for Panels A, C, and D are at the firm-year level, while Panel B is at the patent-firm-year level. All regressions incorporate firm-fixed and year-fixed effects. Robust standard errors are provided in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Period	t	t+1	t+2	t	t+1	t+2
	(1)	(2)	(3)	(4)	(5)	(6)
	IFR-Related Robot Exposure					
	Panel A: Innovation Similarity					
Model	IV					
Variable	Tech Similarity			BERT Similarity		
Robot Exposure B	-0.011**	-0.010**	-0.011**	-0.029***	-0.034***	-0.035***
	(0.004)	(0.004)	(0.005)	(0.010)	(0.010)	(0.011)
Observations	20,815	18,090	15,909	17,915	15,627	13,794
	Panel B: AI-Directed Innovation					
Model	IV					
Variable	AI Distance			AI Closeness		
Robot Exposure B	-0.217***	-0.127***	-0.112***	0.017***	0.018***	0.017***
	(0.027)	(0.027)	(0.028)	(0.001)	(0.001)	(0.001)
Observations	661,018	607,772	561,545	661,018	607,772	561,545
	Panel C: Innovation Input					
Model	IV					
Variable	R&D/Total Assets			R&D/Sales		
Robot Exposure B	0.011***	0.010***	0.008**	0.037***	0.032***	0.034***
	(0.004)	(0.004)	(0.004)	(0.008)	(0.008)	(0.009)
Observations	20,815	18,090	15,909	20,815	18,090	15,909
	Panel D: Innovation Output					
Model	Possion			IV		
Variable	Citation-Deflated Patent Number			Citation-Deflated Patent Value		

Table A21: Robustness Checks of Alternative Robot Exposure Measures: Filter Time Fixed Effect

This table presents the results of robustness checks, using alternative robot exposure measures to examine the impact of robot exposure on firms' innovation. We regress robot exposure on time fixed effects and denote the residuals from this regression as *Robot Exposure C*. Panels A, B, and C display the second-stage robustness results of the 2SLS estimation, with Panel A focusing on firms' innovation similarity, including *Tech Similarity* and *BERT Similarity*, Panel B focusing on AI-directed innovation at the patent-firm-year level, including *AI Distance* and *AI Closeness*, and Panel C examining the effects on firms' innovation inputs, including $R\&D/Total\ Assets$ and $R\&D/Sales$. Panel D presents the Poisson regression robustness results for innovation outputs, including *Patent Number* and *Citation-Deflated Patent Number*. Columns (1)–(3) report the coefficients of robot exposure on the first independent variable for years t, $t+1$, and $t+2$, respectively. Columns (4)–(6) present the coefficients for the second independent variable for the same 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. The regression data for Panels A, C, and D are at the firm-year level, while Panel B is at the patent-firm-year level. All regressions incorporate firm-fixed and year-fixed effects. Robust standard errors are provided in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Period	t	t+1	t+2	t	t+1	t+2
	(1)	(2)	(3)	(4)	(5)	(6)
	Robot Exposure with Time Fixed Effects Filtered					
	Panel A: Innovation Similarity					
Model	IV					
Variable	Tech Similarity			BERT Similarity		
Robot Exposure C	-0.123**	-0.139*	-0.171*	-0.303**	-0.471**	-0.529*
	(0.061)	(0.080)	(0.103)	(0.151)	(0.236)	(0.302)
Observations	20,815	18,090	15,909	17,915	15,627	13,794
	Panel B: AI-Directed Innovation					
Model	IV					
Variable	AI Distance			AI Closeness		
Robot Exposure C	-1.437***	-0.831***	-0.796***	0.113***	0.120***	0.124***
	(0.177)	(0.177)	(0.200)	(0.005)	(0.005)	(0.006)
Observations	661,018	607,772	561,545	661,018	607,772	561,545
	Panel C: Innovation Input					
Model	IV					
Variable	R&D/Total Assets			R&D/Sales		
Robot Exposure C	0.127**	0.135**	0.119*	0.412***	0.447**	0.514**
	(0.053)	(0.066)	(0.071)	(0.152)	(0.193)	(0.249)
Observations	20,815	18,090	15,909	20,815	18,090	15,909
	Panel D: Innovation Output					
Model	Possion			IV		
Variable	Citation-Deflated Patent Number			Citation-Deflated Patent Value		
Robot Exposure C	-0.057**	-0.027	-0.018	-0.656**	-0.765*	-1.048*
	(0.023)	(0.025)	(0.024)	(0.313)	(0.430)	(0.587)

Table A22: Robustness Checks of Alternative Robot Exposure Measures: Rolling Weighted Average

This table presents the results of robustness checks, using alternative robot exposure measures to examine the impact of robot exposure on firms' innovation. 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 year t, compute the rolling weighted average, and then standardize the measure to have a mean of zero and a standard deviation of one, denoted as *Robot Exposure D*. Panels A, B, and C display the second-stage robustness results of the 2SLS estimation, with Panel A focusing on firms' innovation similarity, including *Tech Similarity* and *BERT Similarity*, Panel B focusing on AI-directed innovation at the patent-firm-year level, including *AI Distance* and *AI Closeness*, and Panel C examining the effects on firms' innovation inputs, including $R\&D/Total\ Assets$ and $R\&D/Sales$. Panel D presents the Poisson regression robustness results for innovation outputs, including *Patent Number* and *Citation-Deflated Patent Number*. Columns (1)–(3) report the coefficients of robot exposure on the first independent variable for years t, $t+1$, and $t+2$, respectively. Columns (4)–(6) present the coefficients for the second independent variable for the same 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. The regression data for Panels A, C, and D are at the firm-year level, while Panel B is at the patent-firm-year level. All regressions incorporate firm-fixed and year-fixed effects. Robust standard errors are provided in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Period	t	t+1	t+2	t	t+1	t+2
	(1)	(2)	(3)	(4)	(5)	(6)
	Five-Year Rolling Weighted Average Robot Exposure					
	Panel A: Innovation Similarity					
Model	IV					
Variable	Tech Similarity			BERT Similarity		
Robot Exposure D	-0.092**	-0.095*	-0.337**	-0.261**	-0.344**	-0.339**
	(0.042)	(0.048)	(0.137)	(0.115)	(0.136)	(0.138)
Observations	20,815	18,090	13,794	17,915	15,627	13,794
	Panel B: AI-Directed Innovation					
Model	IV					
Variable	AI Distance			AI Closeness		
Robot Exposure D	-1.057***	-0.647***	-0.657***	0.083***	0.093***	0.103***
	(0.130)	(0.138)	(0.165)	(0.004)	(0.004)	(0.005)
Observations	661,018	607,772	561,545	661,018	607,772	561,545
	Panel C: Innovation Input					
Model	IV					
Variable	R&D/Total Assets			R&D/Sales		
Robot Exposure D	0.095***	0.092**	0.077*	0.309***	0.304***	0.334***
	(0.035)	(0.039)	(0.041)	(0.094)	(0.107)	(0.121)
Observations	20,815	18,090	15,909	20,815	18,090	15,909
	Panel D: Innovation Output					
Model	Possion			IV		
Variable	Citation-Deflated Patent Number			Citation-Deflated Patent Value		
Robot Exposure D	-0.044	-0.017	-0.010	-0.492**	-0.521*	-0.682**

Table A23: The impact of the Supreme Court Case Alice on AI Patent Applications

This table examines the impact of the Supreme Court case Alice on AI patent applications. We use two measures to assess the effect on AI patent filings: $Dummy(Grant)$ is a dummy variable, which takes the value of 1 if the patent is ultimately granted, and 0 otherwise, and $Ln(Grant\ Lag)$ is the logarithm of the number of days between the application and the approval date, after adding 1. We also include the interaction term $Dummy(AI\ Patent) \times Dummy(After\ 2014)$ to test the approval difference between AI patents and non-AI patents before and after the Alice decision, specifically after 2014. $Dummy(AI\ Patent)$ is defined as 1 for AI patents and 0 for non-AI patents, while $Dummy(After\ 2014)$ is defined as 1 for patents filed in 2014 and thereafter, and 0 otherwise. Columns (1)-(3) use the OLS method, while column (4) employs the Poisson method. The sample period spans the years 2004–2019, and the regression data is at the patent-firm-year level. Error terms are adjusted for heteroskedasticity and clustered at the firm level. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Level	Patent-Firm-Year			
Period	t			
Variable	Dummy(Grant)		Ln(Grant Lag)	Grant Lag
Model	OLS	OLS	OLS	Possion
	(1)	(2)	(3)	(4)
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)			
Firm FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	Yes	Yes
Observations	3,340,083	3,216,325	659,190	659,185

Table A24: Alice Decision's Impact on Innovation Shift, Input, and Output

This table presents the results of testing the impact of the Alice decision on the main results of the paper. In the regressions, we include the interaction term between *Robot Exposure* and *Dummy*$(After\ 2014)$ to examine the effect of the Alice decision, specifically after 2014, on innovation shift, input, and output. *Robot Exposure* quantifies the degree of robot exposure faced by the firm in a given year, standardized to have a mean of zero and a standard deviation of one. *Dummy*$(After\ 2014)$ is defined as 1 for patents filed in 2014 and thereafter, and 0 otherwise. Panels A, B, and C display the second-stage robustness results of the 2SLS estimation, with Panel A focusing on firms' innovation similarity, including *Tech Similarity* and *BERT Similarity*, Panel B focusing on AI-directed innovation, including *AI Distance* and *AI Closeness*, and Panel C examining the effects on firms' innovation inputs, including *R&D*$/$*Total Assets* and *R&D*$/$*Sales*. Panel D presents the Poisson regression robustness results for innovation outputs, including *Patent Number* and *Citation-Deflated Patent Number*. Columns (1)–(3) report the coefficients for the first independent variable for years t, $t+1$, and $t+2$, respectively. Columns (4)–(6) present the coefficients for the second independent variable for the same 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. The regression data for Panels A, C, and D are at the firm-year level, while Panel B is at the patent-firm-year level. All regressions incorporate firm-fixed and year-fixed effects. Robust standard errors are provided in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Period	t	t+1	t+2	t	t+1	t+2
	(1)	(2)	(3)	(4)	(5)	(6)
	Alice Decision's Impact					
	Panel A: Innovation Similarity					
Model	IV					
Variable	Tech Similarity			BERT Similarity		
Robot Exposure × Dummy(After 2014)	-0.114**	-0.118**	-0.118	-0.267**	-0.342**	-0.550*
	(0.050)	(0.058)	(0.085)	(0.112)	(0.164)	(0.298)
Robot Exposure	-0.037	-0.050	-0.130	-0.180	-0.357*	-0.384
	(0.063)	(0.083)	(0.098)	(0.150)	(0.213)	(0.256)
Observations	20,815	18,087	15,904	17,915	15,624	13,789
	Panel B: AI-Directed Innovation					
Model	IV					
Variable	AI Distance			AI Closeness		
Robot Exposure × Dummy(After 2014)	-0.365***	-0.038	0.158	0.122***	0.089***	0.062***
	(0.121)	(0.102)	(0.107)	(0.003)	(0.002)	(0.002)
Robot Exposure	-0.718***	-0.548***	-0.712***	-0.017***	0.007**	0.026***
	(0.109)	(0.122)	(0.175)	(0.003)	(0.003)	(0.005)
Observations	661,018	607,772	561,545	661,018	607,772	561,545
	Panel C: Innovation Input					
Model	IV					
Variable	R&D/Total Assets			R&D/Sales		
Robot Exposure × Dummy(After 2014)	0.029	0.040	0.069	0.238**	0.220*	0.407**
	(0.039)	(0.045)	(0.058)	(0.098)	(0.114)	(0.175)
Robot Exposure	0.154**	0.152**	0.105	0.312**	0.382*	0.330
	(0.060)	(0.073)	(0.068)	(0.152)	(0.197)	(0.213)

(Continued)

Period	t	t+1	t+2	t	t+1	t+2
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel D: Innovation Output					
Model	Poisson			IV		
Variable	Citation-Deflated Patent Number			Citation-Deflated Patent Value		
Robot Exposure × Dummy(After 2014)	-0.189***	-0.190***	-0.198***	-0.852***	-0.947***	-1.507***
	(0.046)	(0.045)	(0.043)	(0.279)	(0.323)	(0.504)
Robot Exposure	-0.013	0.018	0.024	0.125	0.138	0.100
	(0.030)	(0.026)	(0.025)	(0.374)	(0.454)	(0.504)
Observations	16,468	14,158	12,265	20,815	18,087	15,904
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table A25: Robustness Checks of Controlling for the impact of the Supreme Court Case Alice

This table presents the results of robustness checks, controlling for the impact of the Supreme Court case Alice. Specifically, for the patent-firm-year level regressions in the paper, we include the control variable *Dummy*(*AI Patent* & $After$ 2014), which is defined as 1 for AI patents filed in 2014 and thereafter, and 0 otherwise. Panels A and B display the second-stage robustness results of the 2SLS estimation, with Panel A focusing on AI-directed innovation, including *AI Distance* and *AI Closeness*, and Panel B examining the effects on firms' innovation quality, including *Generality* and *Originality*. *Robot Exposure* quantifies the degree of robot exposure faced by the firm in a given year, standardized to have a mean of zero and a standard deviation of one. Columns (1)–(3) report the coefficients of robot exposure on the first independent variable for years t, $t+1$, and $t+2$, respectively. Columns (4)–(6) present the coefficients for the second independent variable for the same 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. The regression data for Panels A and B are at the patent-firm-year level. All regressions incorporate firm-fixed and year-fixed effects. Robust standard errors are provided in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Level	Patent-Firm-Year					
Model	IV					
Period	t	t+1	t+2	t	t+1	t+2
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: AI-Directed Innovation					
Variable	AI Distance			AI Closeness		
Robot Exposure	-1.027***	-0.614***	-0.603***	0.082***	0.088***	0.093***
	(0.123)	(0.123)	(0.139)	(0.003)	(0.003)	(0.004)
Dummy(AI Patent & After 2014)	-4.941***	-5.119***	-5.214***	0.690***	0.691***	0.700***
	(0.076)	(0.077)	(0.067)	(0.004)	(0.005)	(0.004)
Observations	661,018	607,772	561,545	661,018	607,772	561,545
	Panel B: Generality and Originality					
Variable	Generality			Originality		
Robot Exposure	0.040***	0.023***	0.015**	0.036***	0.018***	0.016***
	(0.005)	(0.005)	(0.006)	(0.005)	(0.005)	(0.006)
Dummy(AI Patent & After 2014)	0.003	0.014**	0.020***	0.004	0.015***	0.017***
	(0.006)	(0.006)	(0.005)	(0.006)	(0.006)	(0.005)
Observations	607,979	556,885	512,310	365,009	330,468	301,361
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table A26: Robustness Checks of Other Technology Shocks: Drop Firms in IT Industry

This table presents the results of robustness checks, droping other technology shocks to examine the impact of robot exposure on firms' innovation. We exclude firms belonging to the internet industry. Panels A, B, and C display the second-stage robustness results of the 2SLS estimation, with Panel A focusing on firms' innovation similarity, including *Tech Similarity* and *BERT Similarity*, Panel B focusing on AI-directed innovation, including *AI Distance* and *AI Closeness*, and Panel C examining the effects on firms' innovation inputs, including *R&D/Total Assets* and *R&D/Sales*. Panel D presents the Poisson regression robustness results for innovation outputs, including *Patent Number* and *Citation-Deflated Patent Number*. *Robot Exposure* quantifies the degree of robot exposure faced by the firm in a given year, standardized to have a mean of zero and a standard deviation of one. Columns (1)–(3) report the coefficients of robot exposure on the first independent variable for years t, $t+1$, and $t+2$, respectively. Columns (4)–(6) present the coefficients for the second independent variable for the same 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. The regression data for Panels A, C, and D are at the firm-year level, while Panel B is at the patent-firm-year level. All regressions incorporate firm-fixed and year-fixed effects. Robust standard errors are provided in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Period	t	t+1	t+2	t	t+1	t+2
	(1)	(2)	(3)	(4)	(5)	(6)
	Drop Firms in IT Industry					
	Panel A: Innovation Similarity					
Model	IV					
Variable	Tech Similarity			BERT Similarity		
Robot Exposure	-0.192**	-0.218*	-0.266*	-0.480**	-0.732*	-0.800*
	(0.096)	(0.124)	(0.159)	(0.240)	(0.375)	(0.471)
Observations	19,290	16,824	14,839	16,529	14,467	12,810
	Panel B: AI-Directed Innovation					
Model	IV					
Variable	AI Distance			AI Closeness		
Robot Exposure	-1.151***	-0.654***	-0.549***	0.068***	0.077***	0.081***
	(0.146)	(0.145)	(0.162)	(0.004)	(0.004)	(0.005)
Observations	604,911	557,017	515,660	604,911	557,017	515,660
	Panel C: Innovation Input					
Model	IV					
Variable	R&D/Total Assets			R&D/Sales		
Robot Exposure	0.200**	0.209**	0.185*	0.649***	0.623**	0.737**
	(0.084)	(0.102)	(0.110)	(0.241)	(0.275)	(0.362)
Observations	19,290	16,824	14,839	19,290	16,824	14,839
	Panel D: Innovation Output					
Model	Poisson			IV		
Variable	Citation-Deflated Patent Number			Citation-Deflated Patent Value		
Robot Exposure	-0.020	0.015	0.024	-0.997**	-1.149*	-1.571*
	(0.030)	(0.030)	(0.028)	(0.490)	(0.656)	(0.891)

Table A27: Robustness Checks of Other Technology Shocks: Drop Firms with ES Systems

This table presents the results of robustness checks, droping other technology shocks to examine the impact of robot exposure on firms' innovation. We excluded companies that installed ES systems. Panels A, B, and C display the second-stage robustness results of the 2SLS estimation, with Panel A focusing on firms' innovation similarity, including *Tech Similarity* and *BERT Similarity*, Panel B focusing on AI-directed innovation, including *AI Distance* and *AI Closeness*, and Panel C examining the effects on firms' innovation inputs, including $R\&D/Total\ Assets$ and $R\&D/Sales$. Panel D presents the Poisson regression robustness results for innovation outputs, including *Patent Number* and $Citation\text{-}Deflated\ Patent\ Number$. *Robot Exposure* quantifies the degree of robot exposure faced by the firm in a given year, standardized to have a mean of zero and a standard deviation of one. Columns (1)–(3) report the coefficients of robot exposure on the first independent variable for years t, $t+1$, and $t+2$, respectively. Columns (4)–(6) present the coefficients for the second independent variable for the same 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. The regression data for Panels A, C, and D are at the firm-year level, while Panel B is at the patent-firm-year level. All regressions incorporate firm-fixed and year-fixed effects. Robust standard errors are provided in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Period	t	t+1	t+2	t	t+1	t+2
	(1)	(2)	(3)	(4)	(5)	(6)
	Drop Firms with ES Systems					
	Panel A: Innovation Similarity					
Model	IV					
Variable	Tech Similarity			BERT Similarity		
Robot Exposure	-0.309**	-0.408*	-0.464	-0.671**	-1.118	-1.171
	(0.142)	(0.231)	(0.293)	(0.341)	(0.691)	(0.822)
Observations	18,654	16,140	14,127	15,938	13,842	12,161
	Panel B: AI-Directed Innovation					
Model	IV					
Variable	AI Distance			AI Closeness		
Robot Exposure	-1.864***	-1.512***	-2.274***	0.004	0.018***	0.005
	(0.144)	(0.175)	(0.333)	(0.004)	(0.004)	(0.008)
Observations	374,927	338,412	310,014	374,927	338,412	310,014
	Panel C: Innovation Input					
Model	IV					
Variable	R&D/Total Assets			R&D/Sales		
Robot Exposure	0.201**	0.221	0.167	0.711**	0.863*	1.001*
	(0.097)	(0.135)	(0.135)	(0.301)	(0.462)	(0.606)
Observations	18,654	16,140	14,127	18,654	16,140	14,127
	Panel D: Innovation Output					
Model	Poisson			IV		
Variable	Citation-Deflated Patent Number			Citation-Deflated Patent Value		
Robot Exposure	0.008	0.049	0.044	-1.093*	-1.405	-1.952
	(0.034)	(0.034)	(0.035)	(0.571)	(0.918)	(1.303)

Table A28: Robustness Checks of Alternative Samples: Drop Oversea Sales $> 50\%$

This table presents the results of robustness checks, using alternative samples to examine the impact of robot exposure on firms' innovation. We exclude firms whose overseas sales account for more than 50% of total sales. Panels A, B, and C display the second-stage robustness results of the 2SLS estimation, with Panel A focusing on firms' innovation similarity, including *Tech Similarity* and *BERT Similarity*, Panel B focusing on AI-directed innovation, including *AI Distance* and *AI Closeness*, and Panel C examining the effects on firms' innovation inputs, including *R&D/Total Assets* and *R&D/Sales*. Panel D presents the Poisson regression robustness results for innovation outputs, including *Patent Number* and *Citation-Deflated Patent Number*. *Robot Exposure* quantifies the degree of robot exposure faced by the firm in a given year, standardized to have a mean of zero and a standard deviation of one. Columns (1)–(3) report the coefficients of robot exposure on the first independent variable for years t, $t+1$, and $t+2$, respectively. Columns (4)–(6) present the coefficients for the second independent variable for the same 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. The regression data for Panels A, C, and D are at the firm-year level, while Panel B is at the patent-firm-year level. All regressions incorporate firm-fixed and year-fixed effects. Robust standard errors are provided in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Period	t	t+1	t+2	t	t+1	t+2
	(1)	(2)	(3)	(4)	(5)	(6)
	Drop Oversea Sales Ratio $> 50\%$					
	Panel A: Innovation Similarity					
Model	IV					
Variable	Tech Similarity			BERT Similarity		
Robot Exposure	-0.187**	-0.212*	-0.260*	-0.460**	-0.716**	-0.803*
	(0.093)	(0.122)	(0.156)	(0.229)	(0.358)	(0.458)
Observations	20,815	18,090	15,909	17,915	15,627	13,794
	Panel B: AI-Directed Innovation					
Model	IV					
Variable	AI Distance			AI Closeness		
Robot Exposure	-1.004***	-0.589***	-0.555***	0.079***	0.084***	0.087***
	(0.124)	(0.125)	(0.140)	(0.004)	(0.004)	(0.004)
Observations	661,018	607,468	561,097	661,018	607,468	561,097
	Panel C: Innovation Input					
Model	IV					
Variable	R&D/Total Assets			R&D/Sales		
Robot Exposure	0.192**	0.206**	0.181*	0.626***	0.680**	0.780**
	(0.080)	(0.101)	(0.108)	(0.230)	(0.293)	(0.378)
Observations	20,815	18,090	15,909	20,815	18,090	15,909
	Panel D: Innovation Output					
Model	Poisson			IV		
Variable	Citation-Deflated Patent Number			Citation-Deflated Patent Value		
Robot Exposure	-0.026	0.012	0.021	-0.997**	-1.162*	-1.592*

Table A29: Robustness Checks of Alternative Samples: Drop Robot Exposure Top 1%

This table presents the results of robustness checks, using alternative samples to examine the impact of robot exposure on firms' innovation. Given that extreme automation levels in certain firms or industries could affect the generalizability of our baseline results, we exclude the top 1% of firms in terms of automation index for each industry-year. Panels A, B, and C display the second-stage robustness results of the 2SLS estimation, with Panel A focusing on firms' innovation similarity, including *Tech Similarity* and *BERT Similarity*, Panel B focusing on AI-directed innovation, including *AI Distance* and *AI Closeness*, and Panel C examining the effects on firms' innovation inputs, including $R\&D/Total\ Assets$ and $R\&D/Sales$. Panel D presents the Poisson regression robustness results for innovation outputs, including *Patent Number* and $Citation\text{-}Deflated\ Patent\ Number$. *Robot Exposure* quantifies the degree of robot exposure faced by the firm in a given year, standardized to have a mean of zero and a standard deviation of one. Columns (1)–(3) report the coefficients of robot exposure on the first independent variable for years t, $t+1$, and $t+2$, respectively. Columns (4)–(6) present the coefficients for the second independent variable for the same 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. The regression data for Panels A, C, and D are at the firm-year level, while Panel B is at the patent-firm-year level. All regressions incorporate firm-fixed and year-fixed effects. Robust standard errors are provided in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Period	t	t+1	t+2	t	t+1	t+2
	(1)	(2)	(3)	(4)	(5)	(6)
	Drop Top 1% Robot Exposure within Each Industry-Year					
	Panel A: Innovation Similarity					
Model	IV					
Variable	Tech Similarity			BERT Similarity		
Robot Exposure	-0.190**	-0.213*	-0.262	-0.467**	-0.723**	-0.813*
	(0.095)	(0.123)	(0.159)	(0.235)	(0.365)	(0.470)
Observations	20,805	18,080	15,905	17,905	15,617	13,790
	Panel B: AI-Directed Innovation					
Model	IV					
Variable	AI Distance			AI Closeness		
Robot Exposure	-0.983***	-0.580***	-0.534***	0.078***	0.084***	0.086***
	(0.124)	(0.125)	(0.140)	(0.004)	(0.004)	(0.004)
Observations	655,102	602,083	556,188	655,102	602,082	556,187
	Panel C: Innovation Input					
Model	IV					
Variable	R&D/Total Assets			R&D/Sales		
Robot Exposure	0.196**	0.208**	0.183*	0.636***	0.686**	0.787**
	(0.082)	(0.102)	(0.110)	(0.236)	(0.298)	(0.386)
Observations	20,805	18,080	15,905	20,805	18,080	15,905
	Panel D: Innovation Output					
Model	Poisson			IV		
Variable	Citation-Deflated Patent Number			Citation-Deflated Patent Value		
Robot Exposure	-0.026	0.012	0.021	-1.015**	-1.174*	-1.616*

Table A30: Robustness Checks of Firm-Level Results

This table presents the results of robustness checks, shows the firm-year-level results of AI-directed innovation. We use the citation-weighted measure to calculate the firm-level variables. Panels A and B display the second-stage robustness results of the 2SLS estimation, with Panel A focusing on AI-directed innovation, including *AI Distance* and *AI Closeness*, and Panel B examining the effects on firms' innovation quality, including *Generality* and *Originality*. *Robot Exposure* quantifies the degree of robot exposure faced by the firm in a given year, standardized to have a mean of zero and a standard deviation of one. Columns (1)–(3) report the coefficients of robot exposure on the first independent variable for years t, $t+1$, and $t+2$, respectively. Columns (4)–(6) present the coefficients for the second independent variable for the same 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. The regression data is at the firm-year level. All regressions incorporate firm-fixed and year-fixed effects. Robust standard errors are provided in parentheses. Coefficients marked with *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Period	t	t+1	t+2	t	t+1	t+2
	(1)	(2)	(3)	(4)	(5)	(6)
	Firm-Level Results					
	Panel A: AI-Directed Innovation					
Model	IV					
Variable	AI Distance			AI Closeness		
Robot Exposure	-5.132*	-8.029*	-6.867	0.007**	0.006*	0.008
	(3.098)	(4.501)	(5.023)	(0.003)	(0.003)	(0.005)
Observations	20,815	18,088	15,908	20,815	18,088	15,908
	Panel B: Generality and Originality					
Model	IV					
Variable	Generality			Originality		
Robot Exposure	0.003	0.006*	0.006	0.025**	0.049**	0.060**
	(0.002)	(0.003)	(0.004)	(0.011)	(0.022)	(0.030)
Observations	20,815	18,088	15,908	20,815	18,088	15,908
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Appendix C: Proofs

The first order condition for the instantaneous profit implies that

$$q_M^{\beta}(f)k_M^{-\beta}(s \mid f) = \frac{1}{\chi_M} \quad \forall s \in [0, \theta_f] \tag{A1}$$

$$q_L^{\beta}(f)k_L^{-\beta}(s \mid f) = \frac{w}{\bar{q}\chi_L} \quad \forall s \in [\theta_f, 1] \tag{A2}$$

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

$$w = \bar{q}_L \chi_L^{1-\beta}. \tag{A3}$$

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

$$\Pi_f = \alpha_M \theta_f q_n(f) + \alpha_L(1-\theta_f)q_L(f), \tag{A4}$$

where

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

A.1 Proof

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

$$1 \equiv \int_0^1 \left[\int_{\theta_f}^1 l(s \mid f) ds \right] df = \int_0^1 \frac{q_L(f)}{\chi_L \bar{q}} (1-\theta_f) \left(\frac{\bar{q}_L \chi_L}{w} \right)^{1/\beta} df \tag{A5}$$

The profit function

$$\begin{aligned} \Pi_f &= \frac{\beta}{1-\beta} \theta_f \chi_M^{\frac{1-\beta}{\beta}} q_M(f) + \frac{\beta}{1-\beta} \left(1-\theta_f\right) \left(\frac{\bar{q}_L \chi_L}{w}^{\frac{1-\beta}{\beta}} \right) q_L(f) \\ &= \frac{\beta}{1-\beta} \chi_M^{\frac{1-\beta}{\beta}} \theta_f q_M(f) + \frac{\beta}{1-\beta} \chi_L^{1-\beta} \left(1-\theta_f\right) \cdot q_L(f) \\ &= \alpha_M \theta_f q_M(f) + \alpha_L \left(1-\theta_f\right) q_L(f) \end{aligned} \tag{A6}$$

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

$$\Pi_f = \alpha_M \theta_f q_M(f) + \alpha_L (1-\theta_f) q_L(f) \tag{A7}$$

We obtain the solution

$$\begin{aligned} &\gamma A_M q_M = \max_{z_M} \left\{ \alpha_M \theta_f q_M - C_M\left(z_M\right) q_M + z_M A_M \lambda_M q_M \right\} \\ &\Rightarrow \\ &\gamma A_M = \max_{z_M} \left\{ \alpha_M \theta_f - C_M\left(z_M\right) + z_M A_M \lambda_M \right\} \end{aligned} \tag{A8}$$

Similarly,

$$\gamma A_L = \max_{z_L} \left\{ \alpha_L \left(1-\theta_f\right) - C_L\left(z_L\right) + z_L A_L \lambda_L \right\} \tag{A9}$$

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

$$\begin{cases} C_M'(z_M^*) = A_M \lambda_M, \\ C_L'(z_L^*) = \lambda_L A_L. \end{cases} \tag{A10}$$

A_M, A_L are determined.

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

$$\text{Robot exposure} f = \frac{\int_0^{\theta_f} k_M(s \mid f) ds}{\int_{\theta_f}^1 k_L(s \mid f) ds} = \frac{\chi_M^{1/\beta}}{\chi_L} \frac{q_M(f)}{q_L(f)} \frac{\theta_f}{1-\theta_f} \tag{A11}$$

We examine how automation degree θ_f affects innovation strategy

$$\frac{\partial z_M^*}{\partial \theta_f} \quad \frac{\partial z_L^*}{\partial \theta_f} \tag{A12}$$

Using the envelope theorem:

$$\begin{aligned} \gamma \cdot \frac{\partial A_M}{\partial \theta_f} &= \alpha_M, \\ \gamma \cdot \frac{\partial A_L}{\partial \theta_f} &= -\alpha_L, \\ \frac{\partial z_M^*}{\partial \theta_f} &= \frac{\lambda_M}{C_M''(z_M)} \cdot \frac{\partial A_M}{\partial \theta_f} = \frac{\lambda_M}{C_M''(z_M)} \cdot \frac{\alpha_M}{\gamma} > 0. \end{aligned} \tag{A13}$$

$\lambda_M, \alpha_M, C_M''(z_M)$ 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,

$$\frac{\partial z_L^*}{\partial \theta_f} = \frac{\lambda_L}{C_L''(z_L)} \cdot \frac{\partial A_L}{\partial \theta_f} = -\frac{\lambda_L}{C_L''(z_L)} \cdot \frac{\alpha_L}{\gamma} < 0. \tag{A14}$$