Automation-Induced Innovation Shift*

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

May 2024

Abstract

We study the impact of exposure to robotics and automation on corporate innovation, which informs how innovation begets innovation. Using advanced language models, we document that firms with high automation exposure witness a decline in technology similarity over time and significantly shift innovative activities toward AI-related ones, which intuitively complement automation. The shift is more pronounced for firms with greater AI-related research experience or generate more data. 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 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 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]. We bridge this gap by investigating for the first time how a firm's exposure to automation and robotics shapes its overall effort and direction of innovation.

To this end, we extract information from 10K filings to quantify a firm's automation exposure. Leveraging the instrument induced by the demographic change of European Union to infer causality [Acemoglu and Restrepo, 2022], we unveil the dynamic interplay between automation exposure and innovation direction. We propose two measures to capture and quantify shifts in the innovation trajectory of firms: technology similarity and semantic similarity. The technology similarity assesses changes in a firm's patent portfolio composition, viewing each portfolio as a vector wherein dimensions represent the firm's patent activity across various USPTO classifications. By constructing and comparing annual patent distribution vectors over a five-year horizon, we apply cosine similarity metrics to trace the evolution of a firm's technological focus. For textual analysis, we leverage PatentBert—a derivative of the large language model BERT, specifically adapted for the patent domain—to derive semantic representations of patent texts. This approach enhances our ability to detect nuanced shifts in innovation themes, facilitating more precise year-over-year comparisons of innovation content.

Our empirical investigation foremost suggests that heightened automation exposure precipitates a strategic reorientation of innovation endeavors within firms. Specifically, firms with pronounced exposure to robotics exhibit a marked decrease in technology similarity. The consistent decline in technology similarity among these firms underscores the profound effect of automation and highlights the adaptive recalibration of innovation strategies in response to tech-

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

nological advancements. Our findings are robust to alternative measures of automation exposure proposed by Acemoglu and Restrepo [2019], who use the industry-level robot exposure captured by the industrial robots usage per thousand workers and firms' segment sale to construct firm-level automation exposure.

We further zeros in on Artificial Intelligence (AI), a technology field we identify as a primary domain towards which firms reorient their innovations in response to escalating robot exposure. The underlying rationale is twofold: AI not only enhances the functionality of robots by facilitating their operation but also thrives on the data harvested by these robots, creating a virtuous cycle that propels further innovations in AI.

We first conduct a text analysis to empirically distinguish AI-related patents, isolating patents with abstracts containing AI-centric keywords as per Cockburn et al. [2018], thus ensuring our focus is squarely on AI innovation. To refine our analysis beyond mere keyword presence and address the limitations of a simplistic AI versus non-AI binary, we introduce two metrics: AI distance and AI closeness. These metrics are designed to quantitatively assess the proximity of a given patent to a conceptual "representative" AI patent derived from an aggregation of AI patents to mitigate the influence of outliers. This approach is underpinned by the construction of a patent-to-patent citation network, wherein AI distance quantifies the citation network distance from a given patent to this AI archetype, and AI closeness evaluates the citation pattern similarity of a patent to the representative AI patent (detailed in Section 3.3).

Our findings reveal a striking trend: increased robot exposure is correlated with a notable decrease in patent distance from AI innovations, accompanied by a simultaneous increase in patent closeness to AI-related technologies. This shift is further corroborated by an observable surge in data generation activities among firms with heightened robot exposure, as delineated through our text analysis. This empirical evidence suggests that firms strategically pivot their innovation efforts towards AI as a direct consequence of integrating robotics into their operations.

Giczy et al. [2022] create a seed dataset, distinguishing between AI-related and non-AI patents, and use a machine learning model to classify all US patents with modest accuracy. This 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). Using panel regression with firm and year fixed effect and other standard controls, we find that firms with high robot exposure witness a significant shift of innovation toward all eight AI patents, and the effect is most significant for the KR, PLAN, and HW.

We further explore how this exposure influences overall innovative activities. Theoretically, the effect of robot exposure on a firm's innovation output could swing in either direction, influenced by factors such as the transforming costs associated with transitioning to new technological fields, the inherent challenges or efficiencies of different research areas, and R&D spending. Our findings reveal a dual trend among firms with heightened robot exposure: a significant uptick in R&D investments is observed, contrasted by a slight reduction in patent outputs in the years following increased automation exposure. Specifically, our instrumental variables (IV) regression underscores 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%.

This discrepancy between R&D spending and patent numbers may stem from 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, firms' transition towards AI, particularly in cutting-edge technologies such as deep learning and reinforcement learning, often compels them to achieve breakthroughs in foundational science and technology. Mezzanotti and Simcoe [2023] show that patenting is more closely tied to development rather than research. Therefore, if the increase in R&D is predominantly driven by more exploratory, original, and fundamental work—typically associated with the research component of R&D—it is plausible that firms may be investing more in R&D while patenting less, albeit with potentially higher-quality patents. To empirically validate the focus on research related to fundamental problems, We adopt the methodology proposed by Hall et al. [2001] to assess a patent's originality and generality, and our findings reveal that AI patents exhibit significantly

higher levels of both originality and generality.

To further dissect the impact of transition costs on innovation direction, 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.

Beyond the short-term effects of robot exposure, we examine its long-term implications. The shift toward AI-driven innovation shows lasting impacts, with AI closeness remaining significantly positive after ten periods. Innovation output also reveals sustained effects, 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. We further investigate the persistence of the decline in patent activity following increased robot exposure. Given the expected reduction in adjustment costs as firms transition to AI, we hypothesize a "J-curve" effect, where initial declines in patent output are followed by a subsequent increase. Our findings indicate that the negative impact of robot exposure on cumulative patent numbers diminishes over time, ultimately dissipating as the observation period extends. This supports the notion that the initial decline is transitory, reflecting temporary adjustment costs rather than enduring effects.

To validate the robustness of our core findings regarding the shift in innovation and innovation productivity, we conduct a series of robustness checks. These include alternative constructions of key variables, controlling for Alice decisions and other technological shocks, varying

sample specifications, testing results at different levels, and performing placebo tests, all to ensure the reliability of our conclusions.

Finally, to understand better the nuanced empirical dynamics behind innovation shifts amidst increasing automation, we introduce a dynamic general equilibrium model that accommodates firms with varying degrees of automation exposure and allows for endogenous innovation choices. Central to our model is that 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. Our findings pivot around a critical revelation: the escalation in robotic 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.

Ours study contributes across multiple dimensions of technological progress research. 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] highlighted the stimulative impact of public policy on vaccine-related clinical trials. Similarly, Beraja et al. [2022] illustrated how government security procurement contracts in China have spurred AI technology development by providing government-collected data. Further, research by Acemoglu and Linn [2004] and Costinot et al. [2019] has shown 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 and its resultant market expansions.

Our contribution is twofold: we underscore the role of recent technology adoption, automation, as a catalyst for directed innovation shifts, a dimension not thoroughly explored in prior works. Moreover, we reveal the secondary effect of such shifts: a reduction in patent output within incumbent fields, a phenomenon previously undocumented but is related to the recent discussion of technological obsolescence [Ma, 2023]. This dual focus enriches the dialogue on directed technological progress and sheds light on the broader implications of embracing

automation for innovation trajectories.

Our investigation therefore adds to the emerging discussion on the observed decline in research productivity, a phenomenon underscored by prior studies. For example, Griliches [2009] discusses the diminishing returns of patents relative to research expenditure. Jones [2009] develop a theoretical model to emphasize the role of technology-specific knowledge accumulation and argue a trend towards later ages of first patenting among inventors, growing research team sizes, and a drift towards specialization in innovation. This trend suggests an increasing demand for knowledge and learning for inventors to approach the innovation frontier over time. Echoing these observations, Bloom et al. [2020] employs industry-level case studies to illustrate a significant rise in research efforts juxtaposed against a stark decline in research productivity.

Positioned within these narratives, our study introduces an intuitive microfoundation. 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. By correlating increased R&D expenditure with a reduction in patent output in the context of rising automation exposure, our findings provide specific empirical evidence that enriches the existing conversation. This suggests that the reallocation of innovation efforts, while strategically beneficial in the context of technological evolution, may also contribute to the broader trend of reduced research productivity observed across industries.

Our research enriches the multifaceted discourse on the determinants of corporate innovation, which has seen a burgeoning interest across a spectrum of factors 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 into the myriad factors fueling 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 shed light on the distributional effects of innovation shifts. 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 II presents the data and construction of instrumental variables. Section III presents our baseline results. Sections IV and V show the main empirical results on technology shift, robustness, and heterogeneity analysis. Section VI introduces the model but relegate detailed derivations to the appendix. Section VII concludes.

2 Data, Variables, and Identification Strategy

This section briefly describes the construction of data and our instrumental variables.

2.1 Data

The data for this study primarily originates from the following sources: First, patent-related data is obtained from the United States Patent and Trademark Office (USPTO). USPTO has made available comprehensive data on patents granted from 1976 to 2024 within the United States.[2] Our study uses patent applications and grants to capture a firm's innovation output. Recognizing the considerable time lag between the application for a patent and its eventual grant—often spanning several years—we base our annual patent count on the application date to reflect more timely innovation activities. Patent data is aligned with corresponding firms using

[2]Detailed information can be accessed at [USPTO Patent Data](https://www.patentsview.org/download/).

the linkage methodology developed by Kogan et al. [2017]. Our dataset excludes entities from the finance and utility sectors and non-domestic firms. Table 1 delineates the composition of our sample, with information on the number of U.S. public firms included, the annual percentage of these firms holding at least one patent, and the total patent count attributed to this group. Secondly, the data pertaining to robotics is sourced from two primary channels. At the firm level, the keywords utilized to compute robot exposure are extracted from the 10-K filings of U.S. publicly traded companies. At the industry level, robotic data is obtained from the International Federation of Robotics (IFR). To construct the instrumental variable, we integrate IFR data with employment statistics and output figures derived from the European Union's EU KLEMS Growth and Productivity Accounts, which provide detailed analyses of capital, labor, energy, materials, and service inputs.Lastly, additional firm-level variables, such as R&D/Total Assets, R&D/Sales, Ln(Assets), Leverage, Ln(age), ROA, Tangibility, are sourced from the Compustat database.

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] employ industry-level robot exposure—defined as robot installations per thousand workers using data from the International Federation of Robotics (IFR)—alongside Compustat segment data to estimate firm-specific exposure. Firm-level robot exposure is computed as follows:

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

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

To address the limitation of relying solely on industry-level robot exposure and concentrated segment business, we integrate 10-K textual filings with data on industrial robot installations per thousand workers. This approach provides a more accurate measure of firm-specific 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.

To refine keyword selection, we implemented two key measures. First, we compiled a comprehensive keyword list from relevant 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 supplemented it with definitions from the International Federation of Robotics (IFR) reports[3]. Second, we reviewed and excluded any terms related to the sale of industrial robots to ensure focus on their adoption and integration. The final list of keywords is presented in Table 2.

To measure firms' robot exposure, 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:

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

where Robot Word Count$_{w,i,t}$ represents the occurrences of keyword w for firm i in year t, and Total Word Count$_{i,t}$ is the total number of words in the same report. This normalization allows consistent comparison of robot exposure across firms. Figure 1 illustrates the relationship between robot exposure and the number of robots used per worker in the United States[4] between 2004 and 2019. As the figure shows, the estimated robot exposure is strongly correlated with

[3]The IFR defines a robot as "an actuated mechanism programmable in two or more axes with a degree of autonomy, capable of performing designated tasks within its environment." Industrial robots are "automatically controlled, reprogrammable, multipurpose manipulators suitable for both stationary and mobile use in industrial automation."

[4]See Section 2.3 for the detail of the calculation of the number of robots used per worker in the United States. And since the IFR only reported the robot applications in the United States, 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 United States accounts for more than 90 percent of the robot market share in North America.

the actual robot installations. Then we standardized robot exposure across all U.S. companies that are not missing robot exposure. After that, we limit our observation to companies that do not have missing control variables. Table A2 provides summary statistics on the prevalence and distribution of robot exposure across industries.

As a robustness check, we propose several alternative measures, with results presented in the appendix. First, we implement a two-stage Lasso regression, following Cong et al. [2019], to select words that significantly explain industry-level robot installations. Second, we measure robot exposure using the textual filings from the past five years, rather than a single year. Third, we measure robot exposure as the ratio of robots used to employment, using the data from IFR, EU KLEMS, and Compustat-Segment dataset. All results remain robust across these alternative specifications.

2.3 Instrumental Variables

Firms actively engaged in AI innovation might disproportionately report robotics in their 10-K filings, potentially confounding our results. We employ an instrumental variable (IV) using the ratio of the aging population to the workforce. The IV is adopted in several recent research Acemoglu and Restrepo [2020][5], who argues that robot usage among the European countries closely mirrors the U.S. trends in the following years in the data and partially reflects the technology side of the industries. Acemoglu and Restrepo [2022] demonstrates that the aging ratio is a crucial determinant of robot adoption rates and varies significantly across countries. Influenced by its demographic features, the established pattern of automation adoption in the European Union serves as an external benchmark for U.S. firms. We hypothesize that the impact of the European Union's automation intensity on U.S. firms' innovation strategies is exogenous, at least not driven by factors affecting the U.S. public filings, following the rationale used in studies on the effects of Chinese imports by David et al. [2013] and Bloom et al. [2016]. Besides, the industry-level robot exposures of the E.U. should be exogenous to the U.S.

[5]Our strategy would be compromised if changes in robot usage in other advanced economies are correlated with adverse shocks to U.S. firms. For instance, there may be common shocks affecting the same industries in the United States and Europe, such as import competition or rising wages, which could cause industries to adopt robots. Also, the decline of an industry in the United States may encourage both domestic producers in the United States and their foreign competitors to adopt robots.

innovation decision as justified by Acemoglu and Restrepo [2020].

To construct the IV, first, we measure the robots using the data from the International Federation of Robotics (hereafter, IFR). The IFR provides detailed information on robot installations and stocks from all significant robot producers worldwide, accounting for roughly 90 percent of industrial robots. The dataset is aggregated at the country and industry level from 1993 to 2019 but is only available for most countries after 2004. The dataset covers the United States, Finland, Sweden, Denmark, Italy, and France. Second, we collect the employment data on European Union from level analysis of capital, labor, energy, materials, and services inputs dataset (EU KLEMS), following the same procedure as Acemoglu and Restrepo [2020]. EU KLEMS covers the United States and most of the EU countries. Combining the robot and employment in each year, we construct the instrumental variable [6] following:

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

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.

3 Robot Exposure and Innovation Shift

3.1 General Innovation Shift

Our first objective is to assess the influence of robotics exposure on the innovation trajectory within firms, explicitly investigating the extent to which increased exposure to robotics catalyzes shifts in innovation focus towards AI-oriented domains. Our analysis unfolds in two stages:

[6] Due to the inconsistency of the industry classification between the IFR and EUKLEMS, we follow the procedure by Acemoglu and Restrepo [2020], we manually pin down 15 industries including Agriculture, Forestry, and Fishing, Mining and Quarrying, Electricity, Gas and Water Supply, Construction, Education, Research and Development, and nine industries in manufacturing sectors including Food and Beverages, Textiles (including apparel); Wood and Furniture, Paper and Printing, Plastics and Chemicals Products, glass, ceramics, stone, and mineral products; basic metals and mental products; metal machinery; electrical/electronics; automotive and other vehicles; and miscellaneous manufacturing.

we first explore the broader trends in technology shifts among firms, then narrow our focus to the impact of robotic exposure on the propensity for firms to pivot their innovation efforts towards AI-related fields. Initial findings reveal that firms with heightened exposure to robotics demonstrate a marked shift in their innovation outputs, moving from established to emergent technological domains. This transition is evidenced by a decrease in the similarity of patent portfolios over time, suggesting a significant redirection of innovation focus.

To quantitatively measure the evolution of firms' technological focus, we employ a vectorization approach to represent each firm's patent innovations, leveraging cosine similarity as a metric to gauge technological congruence. This analysis is conducted through three 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 United States Patent and Trademark Office (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 Cooperative Patent Classification (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 introduce a measure of Technology Similarity, defined as follows:

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

The Technology Similarity measure, calculated using the cosine similarity between a firm's technology distribution vectors at time t and $t+5$, is a practical and valuable indicator of a firm's technological focus over five years. Following Bloom et al. [2013], Yang and Zhu [2020], we use the five-year window to measure the technology distribution to smooth time fluctuations. 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 insights into a firm's innovation strategy and offers a tool for predicting and managing technological shifts within firms, making it a crucial component of our research.

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. This measure acts as a proxy for assessing changes in a firm's R&D focus: as firms allocate more resources towards emerging technology fields, deviating from their established domains, we would observe an increase in patents within these new areas at the expense of the incumbent fields. Such a strategic redirection is expected to manifest as a decrease in technology similarity, indicating a broader diversification or shift in technological focus.

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:

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

Here, $Technology\ Similarity_{i,t+j}$ is quantified using the cosine similarity of the firm's patent distributions between year t and $t+5$. To ensure the robustness of our findings, we further analyze the technology similarity over a shorter period, comparing years t and $t+3$. Although this adjustment narrows the scope of our observations—since calculating technology similarity necessitates that firms hold patents from $t-4$ to t—it provides an additional lens through which to assess the influence of robotic technologies on innovation trajectories.

Text Semantic Similarity. To overcome the limitations of the Bag-of-Words model, which treats words in isolation and ignores semantic relationships within patent abstracts, we integrate a more advanced natural language processing (NLP) technique. Specifically, we use a pre-

trained BERT model, PatentBert, fine-tuned on a comprehensive dataset of patent documents Bekamiri et al. [2021]. This approach captures the semantic connections and contextual relevance of words within patent abstracts and claims. PatentBert transforms each patent document into a 768-dimensional continuous vector, providing a dense representation of its semantic content. These vectors are concatenated to create a unified representation for each patent, enabling a nuanced analysis of patent similarity as outlined in Equation 4. Additionally, we adjust the semantic representations using patent citations to account for patent quality.

Our empirical investigation, as detailed in Table 3, yields significant results regarding the effect of robotic exposure on innovation shift captured by the technology similarity.

The first-stage results of our IV regression, in the first column, confirm that weak instrument concerns are unwarranted.The second column highlights the significant impact of robot exposure on technological innovation. Our analysis reveals a substantial shift in the technological distribution over the subsequent years, driven by a rising robot exposure. Specifically, a one standard deviation rise in robot exposure leads to a 0.881 reduction in technology distribution similarity in the first year and a 0.700 reduction in the second year in terms of semantic similarity. These effects are both statistically significant and large in magnitude, particularly when compared to the standard deviation of 0.428. This finding underscores the transformative effect of robot exposure, compelling firms to diversify their technological portfolios. Future analyses will expand on these results by incorporating additional similarity measures, deepening our understanding of how robotic technologies reshape innovation trajectories.

3.2 AI-Oriented Innovation Shift

Building on our findings of a significant orientation towards new technology fields by firms with heightened robot exposure, this section delves deeper to elucidate the specific domains of innovation attraction. Notably, evidence suggests a pronounced pivot towards Artificial Intelligence (AI)-related innovations among these firms. Supporting our empirical observations, a dynamic general equilibrium model detailed in the appendix—which incorporates endogenous decisions for R&D expenditure and innovation—illustrates that firms intensifying their engagement with robotic technologies are naturally inclined to enhance their innovation efforts in areas

that synergize with robotic applications.

The Cooperative Patent Classification (CPC) system, as employed by the USPTO, categorizes technology domains across a spectrum ranging from Human Necessities (A) to Electricity (H), along with a particular category (Y) for marking new technological developments. This system's patents are meticulously assigned to one or more classes and subclasses to reflect their invention scope and technological relevance accurately. However, the broad and evolving nature of AI presents a challenge. The CPC framework does not explicitly delineate AI-specific categories, leading to difficulties in directly mapping patents to AI innovations based on classification codes alone.

We employ a text analysis methodology focused on the patent abstracts to address this challenge and capture the emergence of AI-related innovation among firms with significant robot exposure. This approach allows us to bypass the limitations inherent in the CPC classification by identifying AI-related patents through the prevalence of AI-centric terminology within patent abstracts.

To delineate AI-related patents, we adopt a text analysis protocol, scrutinizing the abstracts of patent documents for the presence of specific AI-related keywords as outlined by Cockburn et al. [2018]. This process identifies a patent as AI-related if its abstract mentions at least one of the predetermined keywords. The selected keywords span three pivotal AI domains: Robotics, Symbolic Systems, and Neural Networks, as detailed in Table 4 following Cockburn et al. [2018]

Our analysis reveals a discernible shift in patent activity towards AI-oriented innovations between 2004 and 2019, underscoring the growing emphasis on AI technologies. To quantify this shift, we aggregate patent abstracts annually and compute the occurrence of each AI keyword within that year's patent pool. The frequency of a keyword w in a given year t is calculated as follows:

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

This frequency metric reflects the prevalence of keyword w among the patents granted in year t, offering insights into the relative emphasis on AI technologies.

To further elucidate the temporal dynamics of AI-related innovation, we assess the change in keyword frequency over the studied period relative to the base year (2004), defining 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, facilitating a clear comparison of AI focus intensity over time.

An analysis of keyword trends within patent documents reveals a significant shift toward AI technologies, particularly in the latter part of the study period. As illustrated in Figure 2, AI-related terms, such as "deep learning," "reinforcement learning," "natural language processing," and "machine learning," became the most prominent by 2019. For example, the frequency of "natural language processing" in 2015 was eight times higher than in 2004, increasing to over 20 times by 2019. Similarly, "machine learning" surged to more than 70 times its 2004 frequency, underscoring the rapid expansion of AI-related innovation. This shift toward AI-focused patents accelerated notably after 2010, as demonstrated by the comparative rise in keyword frequencies between 2010 and 2015.

Complementing the keyword analysis, Figure 3 tracks the evolution of AI patents as a proportion of total U.S. patent filings from 2004 to 2020. This metric, based on the presence of AI keywords or citations of patents embedding these keywords, reveals a sharp rise in AI-related patent activity. The share of AI patents grew from approximately 0.5 percent in 2004 to 2.5 percent by 2020.

These findings underscore a significant and accelerating shift toward AI innovation, reflecting broader transformations in the technological and intellectual property landscapes. The surge in AI-related patents highlights the growing importance of AI across industries and signifies a pivotal shift in R&D focus and patenting strategies.

3.3 AI Distance and Closeness

We first aggregate the AI patents into a singular 'representative' AI patent and construct an annual patent-to-patent citation network. In this network, patents p and p' are considered directly connected if a citation link exists either from p to p' or vice versa. The 'distance'

between two patents is defined as the number of citation links that must be traversed to connect them, with a direct citation constituting a distance of one. This innovative approach allows us to map the intricate relationships between AI and other technologies within the patent landscape.

A patent's "AI distance" is the shortest path to link it to the "representative" AI patent within the citation network. Notably, AI patents are assigned an AI distance of zero, indicating their direct association with AI technology.

We further employ a method based on citation distributions that measure the "closeness" of individual patents to a collective representation of AI technology. The process begins with consolidating all AI-related patents identified via keywords into a singular "representative" 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 $C_{p,t}$:

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

where $||x||_2$ denotes the L_2 norm of vector x, facilitating the normalization of the vectors for a fair comparison.

This measurement of AI Closeness is a nuanced indicator of how closely a patent aligns with the forefront of AI technology, as reflected by its citation practices. Through this innovative

approach, we can discern how firms orient their patenting activities towards AI, offering insights into the broader trend of AI integration within the innovation landscape.

We study whether a firm's robot exposure induces its patents towards AI in terms of the AI closeness and distance using the specification,

$$y_{i,p,t+j} = \alpha_i + \alpha_t + \beta \times Robot\ exposure_{it} + X'_{it}\gamma + \varepsilon_{i,p,t,}, \tag{8}$$

where $y_{i,p,t+j}$ is the logarithm of AI distance plus one or the AI closeness for patent p of firm i at year $t+j$. $Robot\ exposure_{it}$ is the one-period lagged robot exposure of firm i, α_i,α_t are firm and year fixed effect, and X_{it} are the other controlling variables.

Table 5 presents the second-stage results of the IV estimation. Columns 1 and 2 report the IV estimates for the effect of one- and two-year lagged robot exposure on AI distance, while Columns 3 and 4 present the corresponding results for AI closeness. The results indicate that firms with higher robot exposure tend to innovate closer to existing AI patents. Specifically, a one standard deviation increase in one-year lagged robot exposure reduces the AI distance by approximately 0.581, which is statistically significant at the 1% level. Additionally, firms with high robot exposure are more likely to generate patents with citation patterns similar to AI patents. On average, a one standard deviation increase in robot exposure leads to a 0.84 increase in AI closeness in the following year.

3.4 Alternative Measure of the AI-Patents

To address the growing interest of policymakers and researchers, the US Patent and Trademark Office (USPTO) launched a project to compile an artificial intelligence (AI) patent dataset Giczy et al. [2022]. Initially, a seed dataset was created through manual review, distinguishing between AI-related patents and non-AI patents. Subsequently, a machine learning model (Long and Short Term Memory method) was employed to classify all US patents. This 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). However, a significant limitation of this dataset is the relatively low accuracy

of its classification.

In our robustness analysis, we use this dataset to examine whether increased robot exposure influences firms to shift their innovation efforts towards AI. The dataset categorizes patents by AI subclasses, allowing for a detailed investigation. To address concerns regarding classification accuracy, we aggregate patents from each AI subclass into a single "representative" patent for each year. We assess AI 'closeness' and 'distance' using the methodology outlined in Section 3.3.

Our results, presented in Table 6, show that firms with higher robot exposure tend to focus more on AI innovation across all subclasses, as evidenced by increased AI closeness. Panel A indicates that firms with higher robot exposure exhibit significantly higher AI closeness, with all coefficients being economically and statistically significant. Panel B examines AI distance, where we predominantly find negative coefficients, with three subclasses—evolutionary computation, NLP, and hardware—showing significantly negative results. This implies that firms with greater robot exposure are closer to being pioneers within each AI subclass in the citation network.

3.5 Rising Data Generation

This subsection explores a potential mechanism by which increased exposure to robotics drives a shift toward innovation in artificial intelligence (AI) domains. A notable example is Amazon's acquisition of Kiva Systems in March 2012. Amazon, one of the world's largest e-commerce and cloud computing companies, has a leading AI system. Kiva Systems, on the other hand, specialized in developing mobile robotic systems for warehouse automation. The integration of Kiva's robotics technology provided Amazon with access to vast amounts of valuable data, fostering the potential for AI advancements. —- The data generated from Kiva's robotic systems spans several dimensions. First, operational data such as position, speed, acceleration, and path trajectory were collected, enabling Amazon to optimize the movement of robots within warehouses. This operational data helped improve efficiency and reduce errors. Additionally, performance data, including cycle times, output metrics, and accuracy measurements, was gathered to assess and enhance overall system performance.

Moreover, the systems captured environmental data, such as temperature, humidity, vibration, and noise levels, which contributed to maintaining optimal working conditions and minimizing disruption. Energy consumption data, including power usage, power consumption rates, and current flow, was also crucial for managing energy efficiency and reducing costs. Finally, maintenance data, consisting of service life information, fault and alarm records, and maintenance history, allowed for better planning and predictive maintenance, reducing downtime and extending the lifespan of equipment.

These rich data streams provided by the robotics systems not only improved operational efficiencies but also offered substantial insights for training AI models and advancing decision-making capabilities. Drawing on the findings of Beraja et al. [2022], which suggest that government procurements in China significantly bolster AI innovations by granting access to proprietary data, our analysis examines whether firms with substantial robotics exposure similarly gain enhanced access to critical data. This enhanced access could drive the acceleration of AI innovation, leading to breakthroughs in AI technologies and their applications.

Measuring the extent of data accessible to firms poses a considerable challenge. Adopting a methodology akin to that used for assessing robotics exposure, we employ the frequency of terms related to data generation mentioned in 10-K filings as a proxy for data access. The keywords, detailed in Table A4, relate to a broad spectrum 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]. This analysis involves calculating the incidence of these keywords within 10-K annual reports and normalizing them by industry-specific frequencies. The mathematical representation is as follows:

$$\text{Data Generation}_{i,t} = \sum_{w} \frac{\text{Data Word Number}_{w,i,t}}{\text{Data Word Number}_{w,t}}, \tag{9}$$

where $\text{Data Word Number}_{w,i,t}$ denotes the count of data-generating related keywords w for firm i during year t.

Table 7 provides evidence that firms with greater exposure to robotics experience a signif-

icant increase in data generation. Columns 1-3 present the average effects of rising robotic exposure on data generation in the current year and the subsequent two years, while Columns 4-5 explore heterogeneity. Specifically, we divide firms into three size groups: small firms (bottom 30%), medium-sized firms (30%–70%), and large firms (top 30%) in terms of sales and net income.

Our findings indicate that larger firms benefit the most, experiencing the greatest increase in data generation as robotic exposure rises. This outcome aligns with the non-competitive nature of data, where larger firms are better positioned to capitalize on these advancements.

4 Aggregate Effect of Innovation Shift

Amid rising robotization, firms are experiencing a significant shift in innovation dynamics, yet the aggregate impact on innovation remains ambiguous. This segment delves into the collective influence of robot exposure on innovation activities. In theory, the shift in innovation could lead firms to either augment or diminish their innovation output. This variance is contingent upon the transaction costs of moving to new domains, the comparative difficulty or efficiency across research fields, and R&D investment levels. We first documented a marked increase in R&D spending and a reduction in citation-weighted patent output at the individual firm level. Subsequently, we present evidence suggesting that this divergent trend between innovation inputs and outputs can be attributed to more foundational research with greater originality and generality, substantial investments in human capital and the incremental costs associated with AI-related innovations.

4.1 Rising R&D Expenditure and Declining Patent

Patent innovation. Our analysis begins by estimating the effects of robot exposure on firm-level innovation, examining both the outputs, as indicated by patent production, and the inputs, represented by R&D expenditures. Following Cohn et al. [2022], we model the patent grant as a time-varying Poisson process and estimate the underlying arrival rate as:

$$\mathrm{E}\left(\mathrm{y}_{\mathrm{it+j}} \mid \mathbf{X}_{\mathrm{it}}\right) = \exp\left(\text{Robot Exposure}_{it}\beta + \mathbf{X}'_{it}\gamma\right) \tag{10}$$

In this equation, $\mathrm{E}\left(\mathrm{y}_{\mathrm{it+j}} \mid \mathbf{X}_{\mathrm{it}}\right)$ 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 next five years after its application and deflate the citation by the average within each 3-digit CPC class. Robot Exposure$_{it}$ measures the degree of robot exposure by firm i in year t. The vector X_{it} includes other standard firm-level controls, such as total assets (to gauge firm size), firm age, leverage ratio, and return on assets. We incorporate firm-fixed effects to mitigate any bias from time-invariant, firm-specific attributes and year-fixed effects to adjust for global temporal trends. Error terms are adjusted for heteroskedasticity and clustered at the firm level to ensure robustness in our regression analysis.

Our analysis results are presented in Table 8, where the robot exposure metric is standardized for ease of interpretation, with a mean of 0 and a standard deviation of 1. This table outlines the effects of robot exposure on patent generation across the entire sample, including both patent-producing and non-patent-producing firms. Specifically, Columns 1 through 3 display the coefficients for robot exposure in year t on the number of patents in years t, $t+1$, and $t+2$, respectively. We find that increased robot exposure is significantly associated with a decline in patent innovation in years $t+1$ and $t+2$. Each result is both economically meaningful and statistically significant at the 5% level.

To account for patent quality heterogeneity, Columns 4-6 report the effects on citation-adjusted patent outcomes, using citations received in the five years following the patent application. Taken together, our evidence indicates that rising robot exposure is linked to a significant decline in patent production, regardless of whether patent quality adjustments are made.

R&D. We examine the effect of robot exposure on innovation from the perspective of R&D. One possible explanation for the negative impact on patents is that firms shrink their R&D expenditure in response to robot exposure. We examine the impact of robot exposure on the concurrent R&D considering a lag between R&D expenditure and patents. We estimate the following model,

$$y_{i,t} = a_i + a_t + \beta_1 \times \textit{Robot exposure}_{i,t} + \beta_2 \times \textit{R\&D Dummy}_{i,t} + X'_{it}\gamma + \varepsilon_{i,t} \tag{11}$$

where $y_{i,t}$ is firm i's R&D/Total Asset, R&D/Sales in year t, $R\&D\ Dummy_{it}$ is an dummy variable to indicate whether the firm i has R&D expenditure in year t. X_{it} includes the same standard control variables as in equation 10.

Table 9 presents second-stage results of the IV estimates. Columns 1-6 display the coefficients of robot exposure on R&D expenditure, normalized by total assets and total sales, respectively. Contemporary, a one standard deviation increase in robot exposure leads to a 19.2% increase in R&D to total assets and a 62.6% increase in R&D to total sales. These results are both economically and statistically significant at the 5% and 1% levels, respectively.

The rise in R&D expenditure, alongside the decline in patent innovation associated with increased robot exposure, may reflect a shift in innovation efforts. This shift could be attributed to more foundational research with greater originality and generality, higher investments in human capital or increased marginal costs related to AI-driven innovation, as discusssed in detail in the next subsections.

4.2 Generality and Originality

A possible explanation for the contrasting patterns in patent innovation and R&D expenditure is that the shift towards AI-related innovation alters the allocation of resources between research and development. Specifically, firms may focus more on research addressing fundamental problems, which leads to a decline in the quantity of patents but an increase in their originality and generality. The distinction between research and development lies in their role in knowledge creation. According to the National Science Foundation, research is defined as the planned, systematic pursuit of new knowledge or understanding, whereas development refers to the systematic use of research and practical experience to produce new and significantly improved goods, services, or processes. Mezzanotti and Simcoe [2023] show that patenting is more closely tied to development rather than research. Therefore, if the increase in R&D is predominantly driven by more exploratory, original, and fundamental work—typically associated with the research component of R&D—it is plausible that firms may be investing more in R&D while patenting less, albeit with potentially higher-quality patents.

In our case, firms' transition towards AI, particularly in cutting-edge technologies such as deep learning and reinforcement learning, often compels them to achieve breakthroughs in foundational science and technology. AI advancements typically depend on a deep understanding of underlying theories, algorithms, and hardware, and interactions among them. To gain a competitive edge in the AI field, firms must propose original solutions, pushing the boundaries of new algorithms, data structures, and processing methods. Thus, AI patent innovation drives firms to revisit foundational issues, focusing on achieving original breakthroughs in AI theories and methodologies, thereby surpassing existing technological limits. Consequently, firms may increase their investment in R&D, but this compositional shift in R&D activities could explain the decline in patenting.

To empirically validate the focus on research related to fundamental problems, we follow the approach outlined by Hall et al. [2001] to measure a patent's originality and generality:

$$Generality_p = 1 - \sum_{j=1}^{N} s_{pj}^2,\ Originality_p = 1 - \sum_{j=1}^{N} q_{pj}^2 \tag{12}$$

where s_{pj} is share of citations that patent p receives from technology class j within 5 years after patent p is issued, q_{pj} is the share of patents that patent p cites in technology class j.Empirically, we define the technology class at the 3-digit CPC level.

Intuitively, the generality measures how widely a patent was cited by other technology classes. In comparison, the originality measures how widely the patent cites other technology classes. One concern of using citations in the measure is that a patent is more likely to be cited by broader technology classes over time. To mitigate such concern, empirically, we measure the generality of a patent filed at year t using the citation it receives in the following five consecutive years. We also control a patent's filing year by its technology class fixed effect to control for any time-varying technology specificity. We consider the following specification,

$$y_{i,p,t+j} = \alpha_i + \alpha_t + \alpha_{c,t} + \beta \times Robot\ exposure_{i,t-1} + X'_{it-1}\gamma + \varepsilon_{ipt} \tag{13}$$

where $y_{i,p,t+j}$ is the originality or generality measure of patent p filed at year $t+j$ by firm i, α_i is the firm fixed effect, α_t is the year fixed effect, $\alpha_{c,t}$ is the patent filing year fixed effect cross the patent class fixed effect where c is patent p's technology class, X_{it-1} includes standard control variables at the firm level, ε_{ipt} is the error term.

Table 11 presents the results. Columns 1-3 and 4-6 show that higher robot exposure leads to an increase in both generality and originality of innovation in the current year and the subsequent two years. Specifically, a one standard deviation increase in robot exposure results in an approximate increase in patent generality by 0.04, 0.022, and 0.022 in years t, $t+1$, and $t+2$, respectively, which is both economically significant (mean originality is 0.369) and statistically significant at the 1% level. Similarly, a one standard deviation increase in robot exposure leads to an approximate increase in patent originality by 0.036, 0.018, and 0.018 over the same periods.

The above results suggest that during the innovation shift towards AI, firms focus on foundational research with originality and generality, leading to an increase in R&D expenditure but a decline in patent output. This finding also provides new insights into the transformation trend of the United States scientific-industrial complex in the context of AI innovation. As noted in studies such as Arora et al. [2018], a key transformation since the 1980s has been the redirection of resources and attention by many leading firms from more exploratory scientific research to more commercially-oriented projects. Our research, however, identifies that the shift towards foundational technological innovation in AI can, in fact, foster more exploratory scientific research.

4.3 Knowledge Diffusion

Another possible explanation for the rising R&D expenditure and decline in the patent output is a high human capital investment that is specific to AI-related innovations. As Isaac Newton said, "Innovation is just standing on the shoulders of the giants". Learning from history can significantly reduce the cost of human capital specific innovation Yang and Zhu [2020]. We use the AI-related innovation experience before 2004 (the initial year of our sample) to proxy for such shifting cost. If AI-related research experience reduces the cost of AI-related innovations in the future, we expect the shift would be more significant for firms with historical AI research.

We measure a firm's AI research experience based on the AI patents granted between 1993 and 2003. Firms are classified as having AI experience if the number of AI patents they hold is above the median; otherwise, they are considered to have less experience. We then examine

whether firms with prior AI experience shift their research focus more significantly using the following specification:

$$y_{i,p,t} = \alpha_i + \alpha_t + \beta_0 \times Robot\ exposure_{it} + \beta_1 \times AI\ experience_{i,03} \times Robot\ exposure_{it} + X'_{it}\gamma + \varepsilon_{i,p,t}, \tag{14}$$

where $y_{i,p,t}$ is the closeness (distance) of firm i's patent p at year t, α_i is the firm fixed effect, α_t is the year fixed effect, $AI\ experience_{i,03}$ 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, X'_{it} are standard control variables at the firm level.

If prior research experience facilitates the shift toward AI-related innovations, we would expect firms with more AI research experience to shift more significantly toward AI fields in response to rising automation compared to their less experienced counterparts. Columns 1 and 2 display the heterogeneous effects of rising robot exposure on innovation shift, measured by AI-distance and AI-closeness.

Two key findings emerge: First, increased robot exposure leads to a significant decline in AI distance and a rise in AI closeness for both groups of firms. Second, the shifting effect is nearly twice as large for firms with substantial prior AI experience.

4.4 Standing on Jones' Shoulders

Besides the increased investment in human capital, the rising marginal costs associated with AI-driven innovation could also explain the higher R&D expenditure and the reduction in patent output. This subsection examines the relative cost of AI patents to non-AI patents in various dimensions. Jones [2009] develop a theoretical model to emphasize the rising cost of accumulating knowledge for inventors to stand on the giant's shoulders. He constructs several measures to capture the cost of the innovation using the inventor's information.

Intuitively, when innovation requires knowledge accumulation in broader fields, it is more likely to be made by a team with a larger team size Jones [2009], where inventors can complement their skill sets. We measure a patent's team size as the number of inventors. Figure 4 shows the average team size of AI and non-AI patents over time and suggests that the team size of AI patents is consistently more considerable than that of non-AI patents. The gap between

the two groups accelerated after 2015, consistent with the rising shift toward AI innovation in the past decade.

A patent's team size may not capture its labor input or efforts by inventors if inventors can simultaneously participate in several projects. The inventor's labor input may be double counted if she is in several projects. To mitigate such concern, we calculate each patent's labor input that deflates the team size's double counting. For example, if an inventor involves two patents applied in one year, the inventor's labor input will be equally split between the two patents. Figure 5 shows the time trend of the average labor input for AI and non-AI patents over the years. The average labor input in AI patents is consistently higher than in non-AI patents, suggesting a high cost in AI innovations.

Hall et al. [2001], Jones [2009] shows that an inventor's originality should significantly and positively correlate with its innovation cost. Intuitively, inventors must search for and build their new knowledge in broader fields to make a more original innovation. Here, we measure an inventor's originality as follows.

$$Inventor\ Originality_{i,t} = 1 - \sum_{j=1}^{N} s_{ji,t}^2 \tag{15}$$

where $s_{ji,t}$ is the fraction of patents by inventor i that falls into technology j in the past five years. A high originality for an inventor implies that she conducts patent innovation in a wide range of technology classes.

For each patent, we calculate its inventor originality as the average originality of its inventors. Figure 6 presents the average inventor originality for AI and non-AI patents over time. The results show that AI patents consistently require higher inventor originality compared to non-AI patents, although the gap has narrowed in recent years.

4.5 The Persistence of the Impact

In addition to the short-term effects of robot exposure discussed earlier, another key question concerns its long-term implications. Tables A6 and A7 analyze the persistence of the impact of rising robot exposure over the subsequent ten years. The results show that the shift toward AI-

directed innovation has significant long-term effects, with AI Closeness remaining significantly positive even after a lag of ten periods. Additionally, innovation output exhibits notable long-term effects, as indicated by the Citation-Deflated Patent Number, which remains significantly negative at the 10% level after eight lagged periods. In contrast, the impact of robot exposure on other variables is predominantly short-term, extending up to a maximum of four lagged periods.

4.6 J-Curve of Number of Patents

We further examine the persistence of the decline in patent activity following rising robot exposure. Given that the adjustment costs associated with firms transitioning toward AI are expected to diminish over time, we hypothesize the emergence of a "J-curve" pattern, where the number of patents initially decreases but subsequently increases. To test this hypothesis, we calculate the cumulative number of patents over one to five years following the initial impact. The results are presented in Table A8. We measure patent activity using both the Patent Number and the Citation-Deflated Patent Number.

As shown, the negative impact of robot exposure diminishes over time, with the effect shrinking to zero as the observation period extends. This suggests that the initial decline in patent activity is transient, not permanent, supporting the notion that the negative effects are due to temporary adjustment costs rather than lasting impacts.

5 Robustness Checks

In this section, we perform robustness checks on the core findings related to innovation shift and innovation productivity. We employ several approaches, including alternative key variables constructions, controlling the Alice decision and other technological shocks, alternative sample constructions, different level results, and placebo test, to ensure the validity of our results.

5.1 Alternative Key Variables Constructions

Alternative Robot Exposure Measurements. We consider alternatives to measure robot exposure. In the first approach, we filter and reweight keywords using a two-stage regression

process, following Cong et al. [2019]. The first stage employs a LASSO regression for variable selection:

$$y_{jt} = \sum_{w \in W} \beta_w x_{w,j,t} + \varepsilon_{jt} \tag{16}$$

where y_{jt} denotes robot exposure in industry j for year t, defined as the number of industrial robots installed per thousand workers, and W represents the set of keywords. We then identify keywords with non-zero coefficients from the LASSO regression and use them in an OLS regression to refine the coefficient estimates. This refined subset of keywords is denoted as W^*, and their corresponding OLS coefficients as $\beta_w, w \in W^*$.

Each firm's robot exposure for year t is then calculated as:

$$\text{Robot Exposure}_{it} = \sum_{w \in W^*} \beta_w x_{i,w,t}, \tag{17}$$

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

Alternatively, instead of using a one-year mention in 10-K filings, we employ a rolling five-year weighted average to measure robot exposure, accounting for factors like depreciation, with an annual discount rate of 0.8.

In the third approach, we use the ratio of robots used to employment. First, we measure the robots using the data from IFR. Second, we collect the employment data on EU KLEMS, following the same procedure as Acemoglu and Restrepo [2020]. To construct the firm-level robot exposure, we combine our industry-level robot exposure dataset with the Compustat-Segment dataset, which documents detailed information on firms' business across industries. The firm-level robot exposure is defined as

$$Robot\ Exposure_{i,t} = \sum_{j=1}^{m} \frac{Sales_{i,j,t}}{Sales_{i,t}} * \frac{Robots_{j,t}}{Workers_{j,t}} \tag{18}$$

where $Sale_{i,j,t}$ is the annual sale of firm i at the segment (industry) j and year t. $Sale_{i,t}$ is the total sale of firm i at year t. $\frac{Robots_{j,t}}{Workers_{j,t}}$ is the industry level robot exposure for industry j at year t.

Tables A9, A10 and A11 present the results of these alternative measurement methods. The magnitude and significance of the core variable coefficients remain consistent with the earlier results.

Alternative Instrumental Variable Construction. We also employ two alternative methods to construct the instrumental variable, in order to assess the robustness of our baseline results.

In the first method, we employ the Bartik instrumental variable approach. The rationale behind this method is that if the robot exposure in U.S. firms and the number of robots per unit of labor across EU countries both vary over time, the changes in these two variables may be driven by common underlying factors. In this approach, we assume that robot usage across EU countries is determined by industry-specific characteristics, which are fixed over time and independent of individual firms' choices. The differing degrees of aging across industries in the early years lead to varying levels of robot exposure across these industries. The annual sales distribution of firms within different industries then determines the extent of robot exposure for each firm. This method combines fixed industry characteristics with the varying sensitivity of firms to industry-level robot exposure over time. Specifically, we replace the time-varying variable in the original IV calculation, which represents the annual proportion of robot usage to employment across EU countries, with a fixed average proportion of robot usage to employment calculated from 2000 to 2004. 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.

The second method involves expanding the robot usage data to include additional EU countries. In our baseline, we only include five countries - Finland, Sweden, Denmark, Italy, and France - in the E.U. to construct the IV due to missing data, which may not be representative. To validate the robustness of these findings, we modify the industry-level instrument in two ways, following the same procedure in Acemoglu and Restrepo [2020]: (1) by including Germany in the five-country average, and (2) by expanding the average to include all nine advanced automation countries in Europe. These nine countries are Denmark, Finland, France, Germany, Italy, Norway, Spain, Sweden, and the United Kingdom.[7] These countries were selected due to

[7]Since data for Norway are unavailable in the EU KLEMS, we impute the Norwegian distribution using em-

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], the robot adoption rates in the five European countries used in the baseline regression lead those of the United States and follow a similar trajectory, while Germany's adoption rate is significantly ahead of the five-country group. Robot adoption rates in Norway, Spain, and the United Kingdom are comparable to those in the United States. Thus, our baseline analysis employs the five European countries' robot adoption rate as the instrumental variable. However, given that the nine countries represent 41% of the global industrial robot market, they also reflect advanced global robot adoption levels. In the robustness checks, we utilize both the five European countries plus Germany and the full set of nine European countries as alternative instruments.

The results, presented in Tables A12, A13, and A14, confirm that none of the alternative constructions affect the conclusions drawn from the baseline analysis.

5.2 The Impact 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 $Dummy(After\ 2014)$

ployment data from Denmark, Finland, and Sweden.

equals 1 for patents filed in or after 2014.

Table A15 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 A16 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 A17 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.

exclude

5.3 Controlling for Other Technological Shocks

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 internet 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) state-

ment, 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 A18 and Table A19, demonstrate that our primary conclusions remain robust after these adjustments.

5.4 Alternative Sample Constructions

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 A20 and Table A21, which remain consistent with the baseline regression results.

5.5 Firm-Level Results of AI-Directed Innovation

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 section, in order to maintain consistency with other results, we use a citation-weighted measure to calculate firm-level variables. Table A22 presents these results, which are largely consistent with the findings in the original study.

5.6 Placebo Test

Following the placebo test methodology of Acemoglu and Restrepo [2020], we aim to avoid the influence of pre-existing trends, such as the possibility that the innovation shift occurred before firms adopted industrial robots. We test the impact of robot exposure during the 2004–2019 period on firms' innovation shift and innovation productivity between 1988 and 2003. The results in Table A23 show that the regression coefficients are not statistically significant, suggesting that, prior to 2004, there was no clear trend—either positive or negative—of innovation among firms with high robot adoption. This confirms that no pre-existing trends have influenced our baseline regression results.

6 Model

This section develops a dynamic equilibrium model to rationalize the observed innovation shift in response to rising robot exposure. All proofs are put in the appendix 7.

6.1 Production

Consider an economy with a continuous time $t \in [0, \infty)$ and a continuum of firms indexed by $f \in [0, 1]$. Each firm operates a continuum of tasks $s \in [0, 1]$. For a firm f, part of the tasks are automated, and the rest are not automated. Suppose tasks in $[0, \theta_f]$ are automated while those in $[\theta_f, 1]$ are non-automated. Acemoglu and Restrepo [2019] use this setup with a continuum of tasks to model the degree of automation and examine the impact of automation on labor and wage inequality. We follow this setup to model firm's automation. Finally, one unit of labor is supplied without elasticity in the market.

Firms make two types of innovations, one type is to improve the productivity of automated tasks while another is to lift up the productivity of the non-automated tasks.

For the tasks in $[0, \theta_f]$, firms make innovations, patent the innovations, and use the patented technologies to improve the productivity of the tasks. The innovations associated with the

automated tasks are combined with the machines to produce the final product:

$$\frac{1}{1-\beta} q_M(f) k_M^{1-\beta}(s \mid f), \tag{20}$$

where $q_M(f)$ is the technology used in automatized tasks, and $k_M(s \mid f)$ is the units of machine used in task s by firm f. To produce one unit of machine, $\frac{1}{\chi_M}$ units of final good are used.

For the non-automatized tasks, firm f combines the innovation adoptable in these lines and labors to produce its final products with

$$\frac{1}{1-\beta} q_L(f) k_L^{1-\beta}(s \mid f), \tag{21}$$

where $q_L(f)$ is the technology used in non-automatized tasks, and $k_L(s|f)$ is produced by labor. Each unit of $k_L(s|f)$ can be produced by $\frac{1}{\chi_L \bar{q}_L}$, where $\bar{q}_L = \int_0^1 (1-\theta_f) q_L(f) df$. We normalize the production of $k_K(s|f)$ by $\bar{q}_L$ to guarantee that the stationary equilibrium is of linear Akcigit and Kerr [2018].

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

$$\begin{aligned}\Pi_f(q_M, q_L) = \max_{k_M(s|f), k_L(s|f)} & \frac{1}{1-\beta} \int_0^{\theta_f} q_M^{\beta}(f) k_M^{1-\beta}(s \mid f)\, ds - \frac{1}{\chi_M} \int_0^{\theta_f} k_M(s \mid f) ds \\ & + \frac{1}{1-\beta} \int_{\theta_f}^1 q_L^{\beta}(f) k_L^{1-\beta}(s \mid f) ds - \int_{\theta_f}^1 \frac{w k_L(s \mid f)}{\bar{q} \chi_L} ds,\end{aligned} \tag{22}$$

where w is the wage for labor which is determined by labor market clear.

Labor market's clear implies $w = \bar{q}_L \chi_L^{1-\beta}$. Firm's first order conditions imply that it's profit takes a linear form with respect to $q_M(f)$ and $q_L(f)$,

$$\Pi_f = \alpha_M \theta_f q_n(f) + \alpha_L (1-\theta_f) q_L(f), \tag{23}$$

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

6.2 Innovation

Innovations in automation tasks. For the innovation process, we follow a similar setup by Akcigit and Kerr [2018] where they model innovation as an arrival process. Specifically, for the technology used in automated tasks, firm f can invest in R&D to promote the arrival of the new technology. The R&D expenditure associated with the arrival rate z_M is

$$C_M(z_M, q_M) = C_M(z_M) q_M \tag{24}$$

with properties that $C_M(0) = 0, C_M'(0) = 0, C_M''(0) > 0$ and that there is an upper bound $\bar{z}$ such that $C_M(\bar{z}) = \infty$.

In this subsection, without confusion, we relabel $q_M(f)$ and $q_L(f)$ as $q_M(t|f)$ and $q_L(t|f)$ to emphasize that technology is time-varying. $q_k(t|f)$ and $q_k(f), k \in \{M, L\}$ are exchangeable without any confusion.

Given the arrival rate z_M, firm f patents one innovation with probability $z_M \Delta t$ and has no innovation with probability $1 - z_M \Delta t$ between t and $t + \Delta t$. The technology of firm f at $t + \Delta t$ is

$$q_M(t + \Delta t \mid f) = \begin{cases} q_M(t \mid f)(1 + \lambda_M) & \text{w.p. } z_M \Delta t \\ q_M(t \mid f) & \text{w.p. } 1 - z_M \Delta t \end{cases} \tag{25}$$

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 $t + \Delta$ is the same as that at t. Otherwise, there is one innovation and the technology is lifted up by λ_M. The probability with two or more innovations is of order $o(\Delta t)$ which is negligible in the context of continuous time.

Innovations in non-automatized tasks. Firm f also takes R&D to improve the technology employed in non-automatized tasks following similar arrival process as in automation tasks. Specifically, the R&D expenditure associated with the arrival rate z_L is

$$C_L(z_L, q_L) = C_L(z_L) q_L \tag{26}$$

with properties that $C_L(0)=0, C_L'(0)=0, C_L''(0)>0$ and that there is an upper bound $\bar{z}$ such that $C_L(\bar{z})=\infty$. Besides, the technology progress between t and $t+\Delta t$ follows

$$q_L(t+\Delta t \mid f)=\begin{cases} q_L(t \mid f)(1+\lambda_L) & \text{w.p } z_L \Delta t \\ q_L(t \mid f) & \text{w.p } 1-z_L \Delta t \end{cases} \tag{27}$$

6.3 R&D Decision

For a firm f, denote the firm's value as V that depends on q_M, q_L, and can be time-varying. Specifically, firm f solves a dynamic optimization problem

$$\begin{aligned} rV-\frac{\partial V}{\partial t}=&\max_{z_M,z_L}\{\Pi_f-C_M(z_M)q_M-C_L(z_L)q_L \\ &+z_M[V(q_M(1+\lambda_M),q_L)-V(q_M,q_L)] \\ &+z_L[V(q_M,q_L(1+\lambda_L))-V(q_M,q_L)]\} \end{aligned} \tag{28}$$

Here, we drop the script f for the notational simplicity but without confusion. The equation is very intuitive. On the right-hand side, firm f obtains the instantaneous cash flow that is the instantaneous profit Π_f from production minus the R&D expenditure, and the future cash flow from the innovation. $z_M[V(q_M(1+\lambda_M),q_L)-V(q_M,q_L)]$ is the net expected value from the innovation in the automation tasks, and $z_L[V(q_M,q_L(1+\lambda_L))-V(q_M,q_L)]$ is the net expected value from the innovation in the non-automation tasks. The left side is the firm's value with a discount rate of r.

We consider the stationary solution to the economy where $\frac{\partial V}{\partial t}=0$ and the Markov Perfect Equilibrium where the value depends on the state variable q_M and q_L. We substitute the profit into the equation [28] and obtain the Bellman equation

$$\begin{aligned} rV=&\max_{z_M,z_L}\{z_M\alpha_L\theta_f q_M+\alpha_L(1-\theta_f)q_L-C_M(z_M)q_M-C_L(z_L)q_L \\ &+z_M[V(q_M(1+\lambda_M),q_L)-V(q_M,q_L)] \\ &+z_L[V(q_M,q_L(1+\lambda_L))-V(q_M q_L)]\} \end{aligned} \tag{29}$$

We present a theorem to state the unique solution of the equation.

PROPOSITION 6.1. *The unique solution to the equation [29] takes*

$$V(q_M, q_L) = A_M q_M + A_L q_L \tag{30}$$

where A_M, A_L are determined by

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

and the arrival rate z_M^*, z_L^* chosen by the firm f is determined from the first oder condition

$$C_M'(z_M^*) = A_M \lambda_M \text{ and } C_L'(z_L^*) = \lambda_L A_L. \tag{31}$$

The next proposition states how the automation affects the innovations in automation and non-automation tasks, specifically, we examine $\frac{\partial z_M^*}{\partial \theta_f}$ and $\frac{\partial z_L^*}{\partial \theta_f}$,

PROPOSITION 6.2. *A high level of automation induces high level of innovations in automation tasks and low level of innovations in non-automation tasks, that is,*

$$\begin{aligned} \frac{\partial z_M^*}{\partial \theta_f} &= \frac{\lambda_M}{C_M''(z_M)} \cdot \frac{\alpha_M}{r} > 0 \\ \frac{\partial z_L^*}{\partial \theta_f} &= -\frac{\lambda_L}{C_L''(z_L)} \cdot \frac{\alpha_L}{r} < 0 \end{aligned} \tag{32}$$

Empirically, we interpret the automation θ_f as the degree of robot exposure to f, and we measure the arrival rate z_M using the patents related to *AI* and robots, and z_L using other patents. From the proposition 6.2, there are several implications for heterogeneous analysis. A reduction in the marginal cost $C_M'(\cdot)$ can boost the innovations in automation tasks. For example, the recent breakthrough in Machine Learning and Deep Learning algorithm could contribute to a reduce in $C_M'(\cdot)$, facilitating the innovation shift.

7 Conclusion

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 robotics 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 run and over the long run, 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.

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, Mirko Draca, and John Van Reenen. Trade induced technical change? the impact of chinese imports on innovation, it and productivity. *The review of economic studies*, 83(1):87–117, 2016.

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.

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.

H David, David Dorn, and Gordon H Hanson. The china syndrome: Local labor market effects of import competition in the united states. *American economic review*, 103(6): 2121–68, 2013.

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.

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.

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.

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 Robots Used per Worker in US

This figure reports the relationship between robot exposure and the number of robots used per worker in the United States from 2004 to 2019, based on annual average values for both variables. To calculate the number of robots used per worker in the United States, we source robot application data from IFR and labor employment data from EU KLEMS.

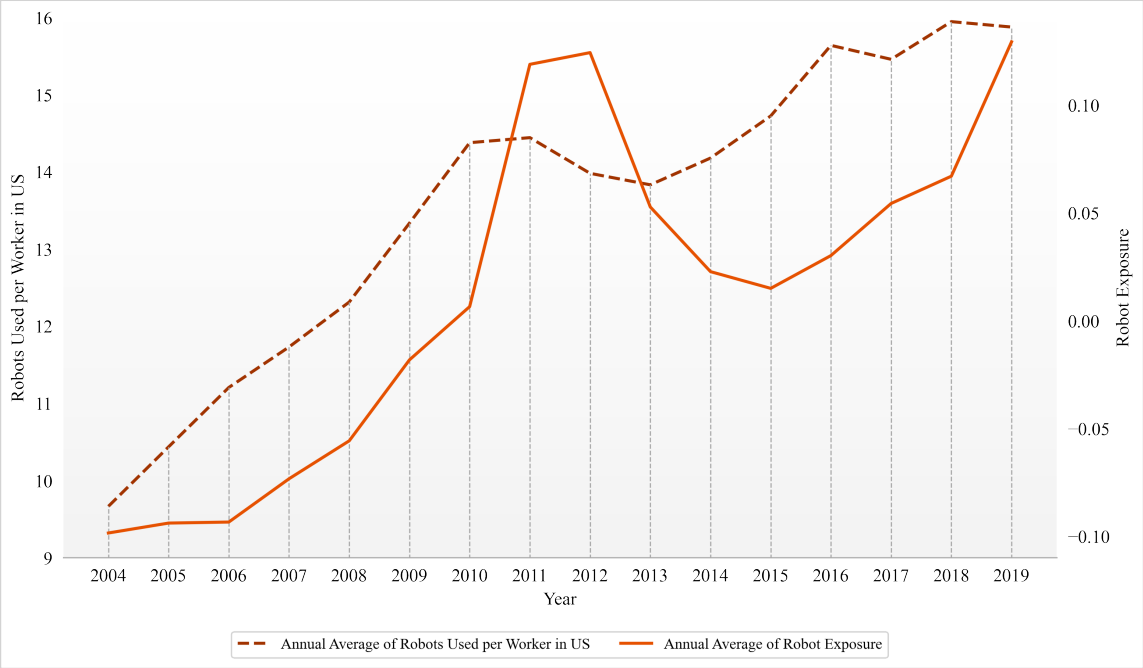


Figure 2: Keywords Frequencies of Patents

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.

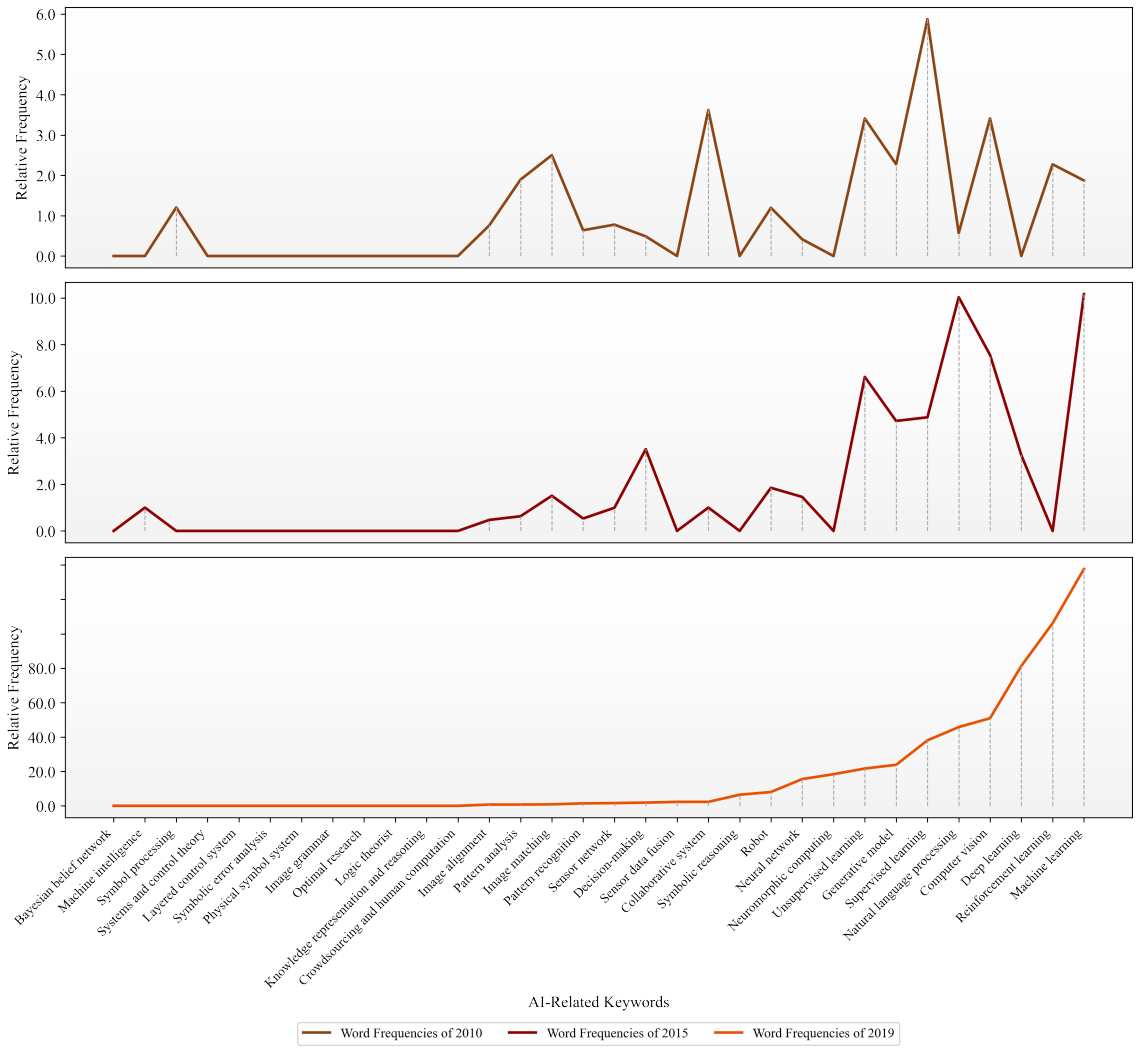


Figure 3: AI Patents in the US from 2004 to 2019

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.

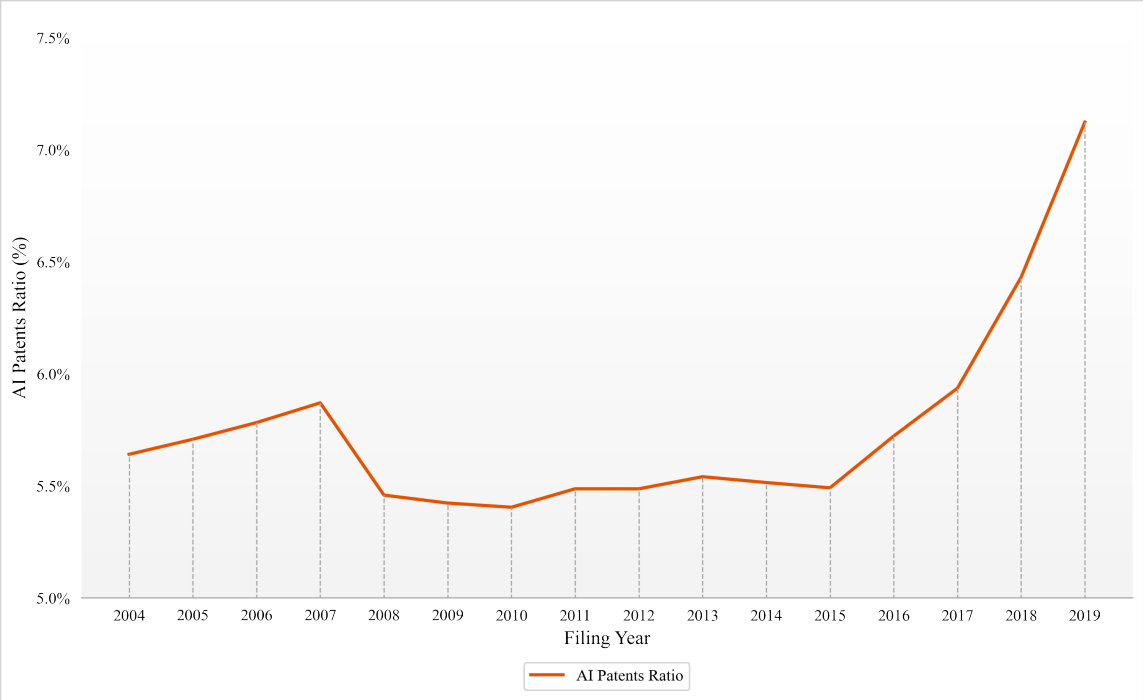


Figure 4: The Team Size for AI and Non-AI Patents

This figure reports the team size for granted patents following Jones [2009], which is the number of inventors involved in the patent, between filing years 2004 and 2019. We divide patents into two categories: AI and non-AI patents. The solid crimson line and the gray dashed line represent the temporal trend in the annual mean team size for AI patents and non-AI patents, respectively.

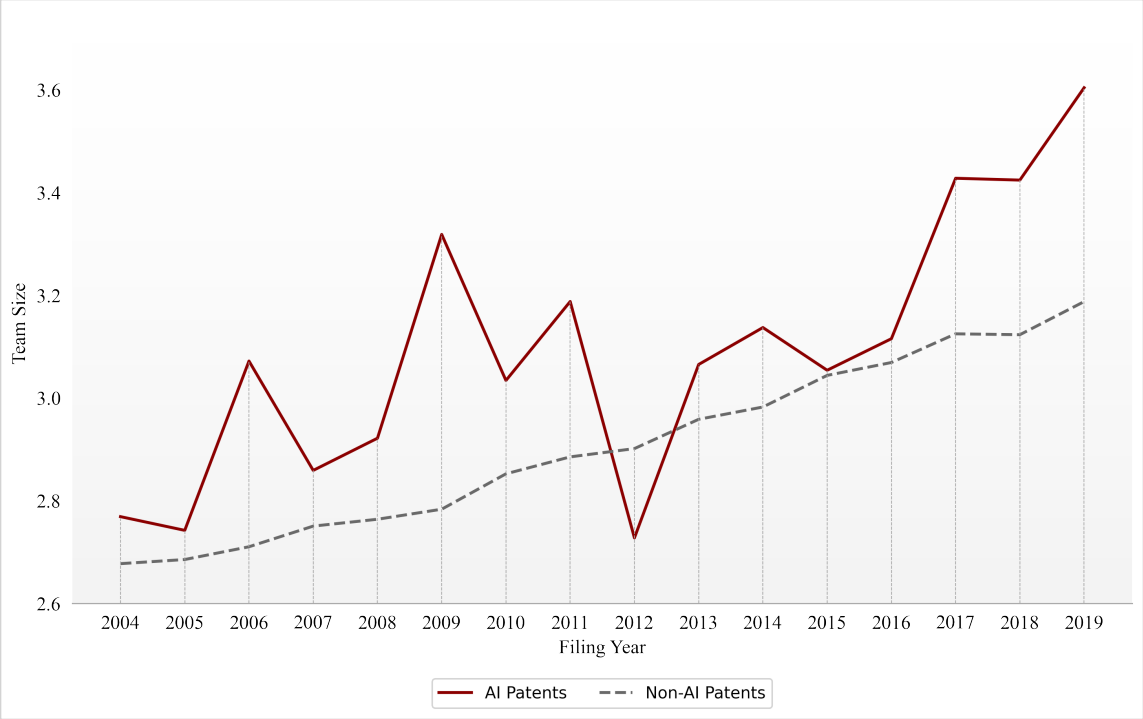


Figure 5: The Labor Input for AI and Non-AI Patents

This figure reports the labor input for granted patents between filing years 2004 and 2019. If an inventor applies for two patents in the same year, we adjust his labor input for each patent as 0.5, and then we aggregate the labor input at the patent level. We divide patents into two categories: AI and non-AI patents. The solid crimson line and the gray dashed line represent the temporal trend in the annual mean labor input for AI patents and non-AI patents, respectively.

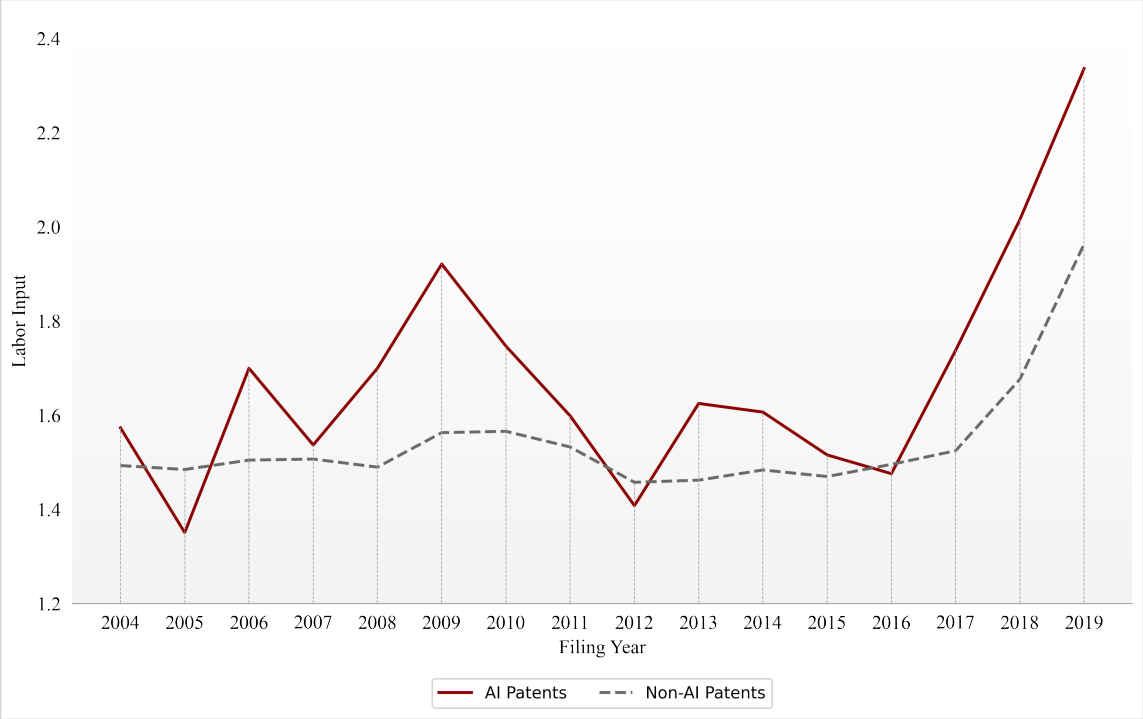


Figure 6: The Inventor Originality for AI and Non-AI Patents

This figure reports the inventor originality for granted patents following Hall et al. [2001] between filing years 2004 and 2019. We measure an inventor's originality as follows.

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

where $s_{ji,t}$ is the fraction of patents by inventor i that falls into technology j in the past five years. Each year, we categorize patents into AI and non-AI groups and calculate the average originality of inventors within each category. The solid crimson line and gray dashed line depict the annual trends in this measure for AI and non-AI patents, respectively.

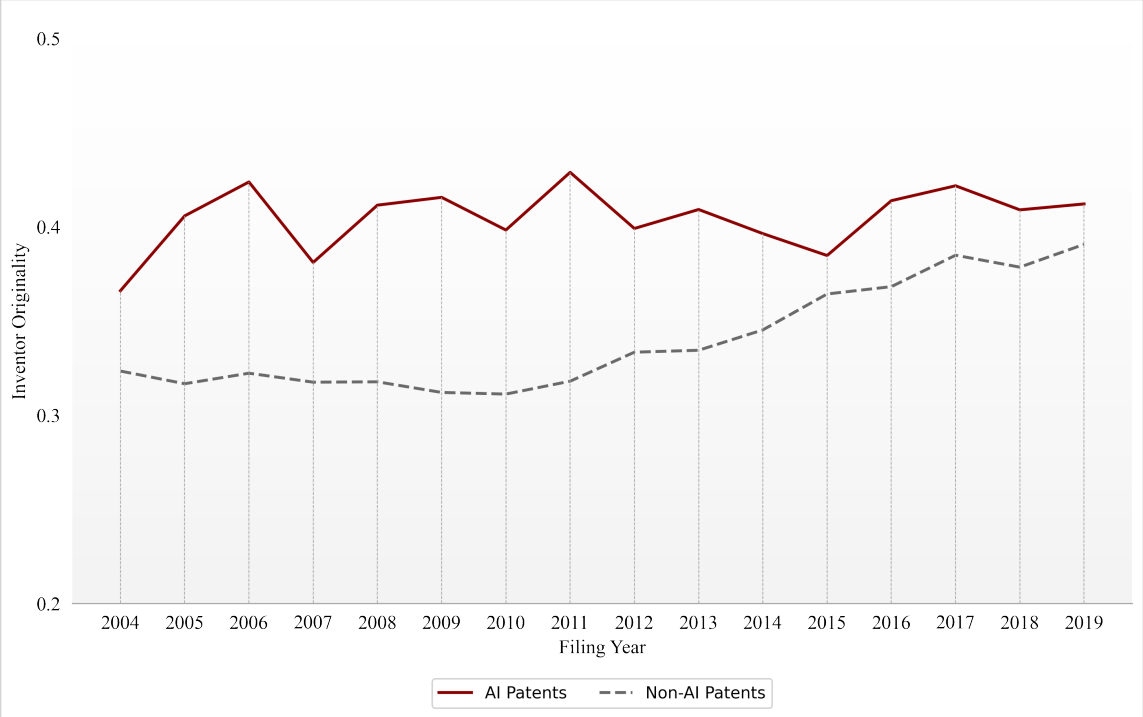


Table 1: Summary of the Match between CompuStat and USPTO

This table provides a comprehensive summary of the match between CompuStat and USPTO data. Column (2) displays the annual count of CompuStat firms that hold at least one patent over the entire sample period. Columns (3) and (4) present the annual count and percentage of firms with at least one patent, respectively, while Column (5) shows the total number of patents granted to CompuStat firms.

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

Table 2: Keyword List for Defining Robot Exposure

This table presents a keyword list used to define firm-level Robot Exposure. The list is developed through a review of relevant literature on robot adoption, with a significant portion sourced from the IFR Reports, which provide detailed definitions of industrial robots and their applications. Column (1) includes the references, while the left columns display the keyword list. In counting keyword frequencies, we do not distinguish between case sensitivity or singular and plural forms of these keywords.

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						
F Statistics (Kleibergen-Paap rk Wald)	14.39						
F Statistics (Cragg-Donald Wald)	16.82						
P-value (Kleibergen-Paap rk LM)	0.000						

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 Distance/AI Closeness

This table presents the second-stage results of the 2SLS estimation, analyzing whether firms adjust their innovations towards AI-related areas in response to increased robot exposure. Specifically, we estimate:

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

where $y_{i,p,t+j}$ is measured using both $\textit{AI Distance}_{i,p,t+j}$ and $\textit{AI Closeness}_{i,p,t+j}$. For each AI field, $\textit{AI Distance}_{i,p,t+j}$ represents the minimum number of steps required for firm i's patent p to reach AI patents within the citation network in year $t+j$. $\textit{AI Closeness}_{i,p,t+j}$ is calculated as the cosine similarity between the citation distribution vectors of firm i's patent p and AI patents in year $t+j$. $\textit{Robot Exposure}_{it}$ 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. Columns (1)–(3) report the coefficients of robot exposure on *AI Distance* for years t, $t+1$, and $t+2$, respectively. Columns (4)–(6) display the coefficients of robot exposure on *AI Closeness* 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	AI Distance			AI Closeness		
Period	t	t+1	t+2	t	t+1	t+2
	(1)	(2)	(3)	(4)	(5)	(6)
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)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	661,018	607,772	561,545	661,018	607,772	561,545

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 Output

This table examines the impact of robot exposure on firms' patent innovation output in subsequent years. We model the patent grant as a time-varying Poisson process and estimate the underlying arrival rate as:

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

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, we consider two measures of patent activity: *Patent Number* refers to the total number of patent applications, while *Citation-Deflated Patent Number* follows the method proposed by Acemoglu et al. [2016]. The latter metric 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. Columns (1)–(3) report the coefficients of robot exposure on *PatentNumber* for years t, $t+1$, and $t+2$, respectively. Columns (4)–(6) display the coefficients of robot exposure on *Citation-Deflated Patent Number* 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					
Variable	Patent Number			Citation-Deflated Patent Number		
Period	t	t+1	t+2	t	t+1	t+2
	(1)	(2)	(3)	(4)	(5)	(6)
Robot Exposure	-0.132***	-0.173***	-0.171**	-0.207***	-0.193**	-0.173**
	(0.040)	(0.065)	(0.082)	(0.069)	(0.077)	(0.087)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	17,400	14,972	13,010	9,951	8,696	7,571

Table 9: 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 10: 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 11: 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: 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: Pairwise correlations

This table tests the validity of the robot exposure variable by comparing its industry-level average with the actual proportion of industrial robots per workforce, as reported by the IFR. Correlations are reported, with coefficients marked *, **, and *** denoting significance at the 10%, 5%, and 1% levels, respectively.

Variable	IFR Robots per Worker	Robot Exposure
IFR Robots per Worker	1.000	
Robot Exposure	0.138**	1.000

Table A4: 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 A5: The relative cost 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:

$$y_{i,p,t} = \alpha_i + \alpha_t + \beta \times Dummy(AI\ patent)_{i,p,t} + X'_{it}\gamma + \varepsilon_{i,p,t},$$

where $y_{i,p,t}$ represents the cost of firm i's patent p in year t. The dependent variable is examined through three 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; and *Inventor Originality* is measured as one minus the sum of the shares of patents filed by the inventor in any given technological class. 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)–(3) report the estimated coefficients of $Dummy(AI\ patent)$ for *Team Size*, *Labor Input*, and *Inventor Originality*, respectively, in year t. Firm-level control variables include firm size (measured by total assets), firm age, leverage ratio, return on assets, and the proportion of tangible assets. The 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	OLS		
Period	t		
Variable	Team Size	Labor Input	Inventor Originality
	(1)	(2)	(3)
Dummy (AI Patent)	0.118***	0.077***	0.041***
	(0.025)	(0.018)	(0.003)
Controls	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Observations	651,475	651,475	621,621

Table A6: 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 ten years. Panel A displays the effects on firms' innovation similarity at the firm-year level, including *Tech Similarity* and *BERT Similarity*. Panel B presents the impact on AI-directed innovation at the patent-firm-year level, including *AI Distance* and *AI Closeness*. *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)–(10) report the estimated coefficients for the effects over the subsequent one to ten 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	t+7	t+8	t+9	t+10
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Panel A: Innovation Similarity									
Level	Firm-Year									
Variable	Tech Similarity									
Robot Exposure	-0.213*	-0.263*	-0.169	-0.035	0.035	-0.011	-0.001	-0.046	0.257	0.728
	(0.123)	(0.158)	(0.137)	(0.131)	(0.128)	(0.136)	(0.210)	(0.369)	(0.940)	(8.459)
Observations	18,088	15,908	13,945	12,198	10,632	9,252	7,916	6,710	5,619	4,592
Variables	BERT Similarity									
Robot Exposure	-0.722**	-0.812*	-0.420	-0.071	0.083	0.029	0.040	0.038	0.075	-0.943
	(0.363)	(0.466)	(0.259)	(0.177)	(0.154)	(0.170)	(0.220)	(0.308)	(2.108)	(3.197)
Observations	15,625	13,793	12,118	10,607	9,259	8,050	6,878	5,841	4,886	3,989
	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	-0.269	-0.623**	0.038	2.764
	(0.124)	(0.140)	(0.155)	(0.176)	(0.188)	(0.196)	(0.227)	(0.316)	(0.468)	(3.470)
Observations	607,772	561,545	517,043	472,706	428,740	391,768	351,068	308,542	259,029	210,866
Variable	AI Closeness									
Robot Exposure	0.084***	0.087***	0.091***	0.102***	0.086***	0.076***	0.062***	0.061***	0.081***	0.912***
	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)	(0.005)	(0.006)	(0.008)	(0.013)	(0.198)
Observations	607,772	561,545	517,043	472,706	428,740	391,768	351,068	308,542	259,029	210,866
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table A7: The impact of Robot Exposure on Innovation Productivity Over the Subsequent Years

This table presents the results of examining the impact of robot exposure on innovation productivity over the subsequent ten years. Panel A reports the second-stage results of the 2SLS estimation, focusing on the effects on firms' innovation inputs, including R&D normalized by total assets (*R&D/Total Assets*) and R&D normalized by total sales (*R&D/Sales*). Panel B presents the Poisson regression results for innovation outputs, including *Patent Number* and *Citation-Deflated Patent Number*. *Patent Number* refers to the total number of patent applications, while *Citation-Deflated Patent Number* follows the method proposed by Acemoglu et al. [2016]. The latter metric adjusts the number of citations each patent received over the subsequent five years, normalizing these citations by the average within each 3-digit CPC class. *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)–(10) report the estimated coefficients over the subsequent one to ten 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 and 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.

Level	Firm-Year									
Period	t+1	t+2	t+3	t+4	t+5	t+6	t+7	t+8	t+9	t+10
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Panel A: Innovation Input									
Model	IV									
Variable	R&D/Total Assets									
Robot Exposure	0.684**	0.784**	0.672**	0.576*	0.307	0.109	-0.078	0.335	-0.413	-11.994
	(0.296)	(0.382)	(0.330)	(0.336)	(0.219)	(0.188)	(0.319)	(0.667)	(1.256)	(131.492)
Observations	18,088	15,908	13,945	12,198	10,632	9,252	7,916	6,710	5,619	4,592
Variable	R&D/Sales									
Robot Exposure	-0.173***	-0.171**	-0.098	-0.106*	-0.101	-0.069	-0.071	-0.123	-0.164	-0.164
	(0.065)	(0.082)	(0.078)	(0.063)	(0.072)	(0.103)	(0.095)	(0.097)	(0.153)	(0.176)
Observations	14,972	13,010	11,224	9,746	8,398	7,145	5,983	4,948	3,995	3,126
	Panel B: Innovation Output									
Model	Possion									
Variable	Patent Number									
Robot Exposure	-0.173***	-0.171**	-0.098	-0.106*	-0.101	-0.069	-0.071	-0.123	-0.164	-0.164
	(0.065)	(0.082)	(0.078)	(0.063)	(0.072)	(0.103)	(0.095)	(0.097)	(0.153)	(0.176)
Observations	14,972	13,010	11,224	9,746	8,398	7,145	5,983	4,948	3,995	3,126
Variable	Citation-Deflated Patent Number									
Robot Exposure	-0.193**	-0.173**	-0.111	-0.141*	-0.173*	-0.201*	-0.240**	-0.222*	-0.178	-0.157
	(0.077)	(0.087)	(0.084)	(0.081)	(0.089)	(0.107)	(0.114)	(0.125)	(0.143)	(0.132)
Observations	8,696	7,571	6,607	5,762	4,933	4,161	3,436	2,796	2,224	1,747
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table A8: The impact of Robot Exposure on the Cumulative Number of Patents

This table examines the impact of robot exposure on the cumulative number of patents over several years. The patent grant process is modeled as a time-varying Poisson process, and we estimate the underlying arrival rate using the following specification:

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

where $\mathrm{E}\left(\mathrm{y_{i,t}} \mid \mathbf{X_{it}}\right)$ represents the expected number of patents accumulated over several years filed by firm i in year t. We consider two measures of patent activity: *Patent Number* refers to the total number of patent applications, while *Citation-Deflated Patent Number* follows the method proposed by Acemoglu et al. [2016]. The latter metric 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. $\textit{Robot Exposure}_{it}$ captures the degree of robot exposure faced by firm i in year t, and is standardized to have a mean of zero and a standard deviation of one. Panel A and Panel B present the results for *Patent Number* and *Citation-Deflated Patent Number*, respectively. Columns (1)–(5) report the Poisson regression results for the cumulative number of patents accumulated over 1 year to 5 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 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				
Period	t				
Cumulative	1 year	2 years	3 years	4 years	5 years
	(1)	(2)	(3)	(4)	(5)
	Panel A: Patent Number				
Variable	Patent Number				
Robot Exposure	-0.132***	-0.110***	-0.074*	-0.011	0.004
	(0.040)	(0.040)	(0.041)	(0.041)	(0.032)
Observations	17,400	15,368	13,456	11,639	9,990
	Panel B: Citation-Deflated Patent Number				
Variable	Citation-Deflated Patent Number				
Robot Exposure	-0.207***	-0.162**	-0.126*	-0.066	-0.037
	(0.069)	(0.067)	(0.070)	(0.064)	(0.043)
Observations	9,951	8,969	8,031	7,107	6,185
Controls	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes

Table A9: 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					
Variable	Patent Number			Citation-Deflated Patent Number		
Robot Exposure A	-0.094***	0.010	0.012	-0.126**	-0.125***	-0.119***
	(0.032)	(0.050)	(0.042)	(0.054)	(0.045)	(0.036)
Observations	17,358	14,951	12,995	9,942	8,692	7,562
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 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 B*. 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 $Citation\text{-}Deflated\ Patent\ Number$. 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 B	-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 B	-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 B	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					
Variable	Patent Number			Citation-Deflated Patent Number		
Robot Exposure B	-0.124***	-0.127**	-0.119*	-0.168***	-0.158**	-0.146*
	(0.045)	(0.059)	(0.066)	(0.061)	(0.071)	(0.076)
Observations	17,358	14,951	12,995	9,942	8,692	7,562
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 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 C*. 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 C	-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 C	-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 C	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					
Variable	Patent Number			Citation-Deflated Patent Number		
Robot Exposure C	-0.042	-0.059	-0.077	-0.130***	-0.143***	-0.161***
	(0.043)	(0.046)	(0.047)	(0.047)	(0.048)	(0.050)
Observations	17,358	14,951	12,995	9,942	8,692	7,562
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table A12: Robustness Checks of Alternative Instrument Variables: Bartik IV

This table presents the results of robustness checks, using Bartik instrument variable to examine the impact of robot exposure on firms' innovation. 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. 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)
	Bartik IV					
	Panel A: Innovation Similarity					
Variable	Tech Similarity			BERT Similarity		
Robot Exposure	-0.258*	-0.277*	-0.210	-0.563*	-0.669	-0.637
	(0.147)	(0.161)	(0.156)	(0.316)	(0.412)	(0.423)
Observations	20,799	18,075	15,895	17,899	15,612	13,780
	Panel B: AI-Directed Innovation					
Variable	AI Distance			AI Closeness		
Robot Exposure	-3.116***	-1.670***	-1.061***	0.126***	0.076***	0.045***
	(0.538)	(0.371)	(0.388)	(0.012)	(0.007)	(0.008)
Observations	660,661	607,117	560,753	660,661	607,117	560,753
	Panel C: Innovation Input					
Variable	R&D/Total Assets			R&D/Sales		
Robot Exposure	0.126	0.130	0.154	0.401*	0.443*	0.435*
	(0.084)	(0.090)	(0.101)	(0.206)	(0.238)	(0.263)
Observations	20,799	18,075	15,895	20,799	18,075	15,895
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 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 A14: 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 A15: 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 A16: 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)
Observations	20,815	18,087	15,904	20,815	18,087	15,904

(Continued)

Period	t	t+1	t+2	t	t+1	t+2
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel D: Innovation Output					
Model	Poisson					
Variable	Patent Number			Citation-Deflated Patent Number		
Robot Exposure × Dummy(After 2014)	-0.277***	-0.418***	-0.405***	-0.436***	-0.413***	-0.388***
	(0.071)	(0.117)	(0.115)	(0.065)	(0.079)	(0.099)
Robot Exposure	0.001	0.001	-0.025	-0.020	-0.035	-0.039
	(0.044)	(0.051)	(0.053)	(0.043)	(0.051)	(0.057)
Observations	17,358	14,949	12,991	9,942	8,691	7,568
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: 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 A18: 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	Possion					
Variable	Patent Number			Citation-Deflated Patent Number		
Robot Exposure	-0.159***	-0.208***	-0.197**	-0.228***	-0.205**	-0.177*
	(0.043)	(0.069)	(0.089)	(0.071)	(0.080)	(0.091)
Observations	16,176	13,940	12,122	9,289	8,114	7,046
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table A19: 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	Possion					
Variable	Patent Number			Citation-Deflated Patent Number		
Robot Exposure	-0.099**	-0.104*	-0.141	-0.204**	-0.181	-0.167
	(0.047)	(0.062)	(0.095)	(0.102)	(0.116)	(0.131)
Observations	15,429	13,217	11,451	8,527	7,421	6,418
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table A20: 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	Possion					
Variable	Patent Number			Citation-Deflated Patent Number		
Robot Exposure	-0.142***	-0.181***	-0.170**	-0.223***	-0.199**	-0.171*
	(0.041)	(0.067)	(0.084)	(0.070)	(0.078)	(0.088)
Observations	17,358	14,951	12,995	9,942	8,692	7,562
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table A21: 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	Possion					
Variable	Patent Number			Citation-Deflated Patent Number		
Robot Exposure	-0.142***	-0.181***	-0.170**	-0.223***	-0.199**	-0.171*
	(0.041)	(0.067)	(0.084)	(0.070)	(0.078)	(0.088)
Observations	17,349	14,943	12,991	9,938	8,690	7,561
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table A22: 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

Table A23: Robustness Checks of Placebo Test

This table presents the results of robustness checks, employing a placebo test to examine the impact of robot exposure on firms' innovation. 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 and B display the second-stage robustness results of the 2SLS estimation, with Panel A focusing on firms' innovation similarity, including *Tech Similarity* and *BERT Similarity*, and Panel B examining the effects on firms' innovation inputs, including *R&D/Total Assets* and *R&D/Sales*. Panel C 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 and 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)
	Placebo Test					
	Panel A: Innovation Similarity					
Model	IV					
Variable	Tech Similarity			BERT Similarity		
Robot Exposure	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: Innovation Input					
Model	IV					
Variable	R&D/Total Assets			R&D/Sales		
Robot Exposure	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
	Panel C: Innovation Output					
Model	Possion					
Variable	Patent Number			Citation-Deflated Patent Number		
Robot Exposure	0.065	0.066	0.068	-0.024	-0.009	0.016
	(0.057)	(0.051)	(0.043)	(0.040)	(0.039)	(0.034)
Observations	14,501	12,996	11,843	8,869	7,990	7,288
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Appendix B: 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

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

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,

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