AI and Productivity: The Role of Innovation*

Viral V. Acharya† Xuehua Shao‡ Zhaoxia Xu§

Abstract

This paper examines an innovation channel through which firms' adoption of artificial intelligence (AI) may affect productivity. Using a novel firm-level measure based on AI product installations, we find that adopters increase patenting and produce patents with more citations and claims, greater originality and generality, and stronger links to existing technologies. Mechanism tests indicate that AI improves research efficiency and reshapes the innovation process through changes in inventor composition and knowledge recombination. Adopters also increase R&D and exhibit higher productivity and firm value. Quantitatively, a Hulten-style approximation implies a 0.61% contribution to aggregate value-added TFP in a representative post-adoption year.

Keywords: Artificial Intelligence, Productivity, Intangible investment, Knowledge recombination, R&D

JEL Classification: D22, D24, E22, L25, O31, O33

*We would like to thank Yakov Amihud, Lerong Cai (discussant), Nikhil Dilip, Yeejin Jang, Byungwook Kim, Youngkyung Kim, Kai Li, Ron Masulis, Gordon Phillips, Robert Seamans, Rik Sen, Elena Simintzi, and Fan Zhang (discussant), seminar participants at NYU Stern Corporate Governance Luncheon Seminar, University of California, Irvine, University of Houston, University of Texas at Austin, and conference participants at the 2026 NBER-SAIF Conference on AI and Financial Markets, the 2026 International Conference on Financial Market and Corporate Finance (ICFMCF), the 2026 Macquarie University Accounting and Finance Conference, and the 2025 Financial Research Network (FIRN) Annual Conference for helpful comments and suggestions.

†New York University Stern School of Business, NBER, ECGI, and CEPR. Email: vva1@stern.nyu.edu

‡University of New South Wales. Email: xuehua.shao@unsw.edu.au

§University of New South Wales. Corresponding author. Email: zhaoxia.xu@unsw.edu.au

1 Introduction

> "*As artificial intelligence evolves, we must remember that its power lies not in replacing human intelligence but in augmenting it. The true potential of AI lies in its ability to amplify human creativity and ingenuity.*"
>
> —Ginni Rometty, *Former CEO of IBM*

As artificial intelligence (AI) is increasingly adopted by firms, a growing literature debates whether and how it affects productivity growth. Task-level evidence documents sizable productivity gains in specific workflows such as customer support (Brynjolfsson, Li, and Raymond, 2025). Yet, realizing productivity gains from AI adoption may require costly adjustments to production processes and complementary organizational investments, generating J-curve dynamics in which short-run productivity losses precede longer-run gains (McElheran et al., 2025). At the aggregate level, task-based approaches yield substantially different estimates of AI's effects on total factor productivity (TFP) (Acemoglu, 2025; Aghion and Bunel, 2024).

Despite growing evidence on AI and productivity, less is known about the mechanisms through which firm-level AI adoption translates into productivity gains. Much of the literature linking AI to productivity has emphasized task-based mechanisms, whereby AI raises productivity by automating or augmenting existing tasks. We instead focus on an innovation channel: AI may reshape the process through which firms innovate, consistent with the view of AI as a method of invention (Cockburn, Henderson, and Stern, 2018). We examine how AI adoption affects firms' innovative activity and productivity, shedding light on the potential role of innovation in AI-driven productivity gains.

The effect of AI on innovation, however, is theoretically ambiguous. On the one hand, AI may facilitate innovation by improving search, accelerating knowledge discovery, and helping firms identify high-value opportunities (Mühlroth and Grottke, 2020; Cockburn, Henderson,

and Stern, 2018; Agrawal, McHale, and Oettl, 2019).[1] Importantly, AI need not directly produce patentable inventions to affect innovation. By augmenting human decision-making and experimentation, AI can improve the processes through which firms discover and develop new products, services, and production methods. On the other hand, the effects of AI on innovation may be constrained by coordination frictions and depend on complementary data, infrastructure, and organizational capabilities (Bresnahan and Trajtenberg, 2017; Cockburn, Henderson, and Stern, 2018; Agrawal, McHale, and Oettl, 2024). AI-generated prior art may also make subsequent inventions more difficult to patent by expanding the volume of prior art, potentially weakening firms' incentives to innovate (Ouellette, Fang, and Ouellette, 2025; Yordy, 2021). Given these competing forces, whether and how AI adoption affects firm innovation remains an empirical question.

One challenge in answering this question is measuring *firm-level* AI adoption and identifying its timing (Raj and Seamans, 2019). Prior studies often proxy for AI investment using AI-related job postings (Babina et al., 2024) or construct measures of AI exposure using Occupational Information Network (O∗NET) data (Eloundou et al., 2024; Acemoglu, 2025; Eisfeldt, Schubert, and Zhang, 2025; Felten, Raj, and Seamans, 2021). While these measures capture important dimensions of firms' AI investment and exposure, they do not directly observe realized adoption and its timing. We address this challenge using newly available historical data on IT product installations to identify firm-level AI adoption and when it occurs. Product installations provide a more direct measure of realized AI adoption, reflecting firms' implementation of AI technologies in their operations. This distinction allows us to examine how realized AI adoption affects subsequent firm innovation and productivity. More broadly, our approach addresses recent calls for more direct measures of firm-level AI adoption (Babina, 2026).

[1] For instance, Microsoft employed Azure Quantum Elements, an AI-powered research platform, to accelerate scientific discovery in chemistry and materials science. In developing a solid-state battery electrolyte, Microsoft researchers used AI models to screen 32.6 million hypothetical materials, reducing the pool to about 500,000. Combining these models with physics-based simulations further narrowed the set to about 150 candidates, after which AI property-prediction filters and expert criteria identified 18 candidates for laboratory testing. The resulting electrolyte was synthesized and validated in less than nine months, a process that would typically take years (McPhee, 2024; Microsoft, 2024). Microsoft subsequently filed a patent application covering the resulting electrolyte compositions (publication No. US2025/0105343 A1). This example illustrates how AI can improve research productivity by reducing the time and experimental resources required to identify viable innovations.

To account for the staggered timing of adoption across firms, we implement a stacked difference-in-differences (DiD) design that compares changes in economic outcomes between firms that first adopt AI in a given year (AI-adopting firms) and firms that have not yet adopted AI (non-adopting firms).[2] For each adoption-year cohort, we include only not-yet-treated firms as clean controls, thereby mitigating bias from treatment effect heterogeneity under staggered adoption (Baker, Larcker, and Wang, 2022).

Using data on granted patents filed by publicly listed U.S. firms to measure innovation outcomes, we find that AI-adopting firms experience subsequent increases in both the quantity and quality of their patents relative to non-adopting firms. AI adoption also influences firms' innovation characteristics along three key dimensions: novelty, direction, and reach. In terms of novelty, AI-adopting firms produce patents that exhibit greater originality and span more technologically distant fields, suggesting that AI enables firms to explore a broader and less familiar technological domains. Regarding innovation direction, we find that AI adoption shifts firms toward more exploitative innovation, which leverages their existing technological capabilities to generate incremental improvements. In relation to innovation reach, we observe that patents produced by AI-adopting firms demonstrate higher generality, implying broader applicability and influence across diverse technology areas. These firms also tend to file patents with more claims, suggesting a potentially broader scope of patent protection.

We next examine the mechanisms through which AI adoption affects innovation. The evidence points to three complementary channels: improvements in research efficiency, shifts in inventor composition, and the recombination of knowledge. Relative to non-adopting firms, AI adopters exhibit higher inventor productivity and greater innovation efficiency. AI adoption also changes inventor composition, with adopters forming larger inventor teams and involving more first-time inventors. Finally, adopters draw on more specialized inventors while assembling teams with a broader range of core technical skills, a pattern consistent with AI facilitating the recombination of specialized knowledge across technological domains. Taken together, these

[2]We use the terms AI-adopting (non-adopting) firms, AI-using (non-using) firms, AI (non-AI) firms, and adopters (non-adopters) interchangeably throughout the paper.

results suggest that AI adoption enhances research productivity and reshapes the innovation process by changing how inventive labor and technological knowledge are organized.

Surprisingly, cross-sectional analyses show that the innovation effects of AI adoption are more pronounced among focused firms, contrary to the view that AI should disproportionately benefit diversified firms operating in more complex information environments. The effects are not stronger among larger or financially unconstrained firms, despite their greater capacity to bear the fixed costs of AI adoption and invest in complementary infrastructure. Moreover, the results are not driven by firms in high-tech sectors, and the innovation effects extend beyond AI-related patents.

Furthermore, AI adopters increase R&D expenditures and experience higher firm value, while physical investment and operating costs do not change significantly relative to non-adopters. Together with the changes in innovation following AI adoption, these patterns suggest that AI adoption is accompanied by greater investment in intangible knowledge capital, rather than physical capital deepening or immediate cost reductions. AI adopters also experience an approximately 8.87% increase in value-added TFP after adoption relative to non-adopters. At the same time, non-adopters exhibit a decline in value-added TFP, consistent with potential spillover or competitive effects. Aggregating these firm-level productivity changes using a first-order Hulten-style aggregation (Hulten, 1978) implies a contribution of approximately 0.61% to aggregate value-added TFP growth in a representative post-adoption year.

A natural concern is that AI adoption may be endogenous, with unobserved factors potentially affecting both adoption and innovation. Our stacked DiD design accounts for time-invariant firm heterogeneity and cohort-specific shocks, and event-study estimates show no evidence of differential pre-trends in innovation before AI adoption. Our results are also robust to placebo tests based on randomly assigned adoption and to entropy balancing that improves comparability between adopters and non-adopters on pre-treatment characteristics. We further address potential endogeneity using a control function approach, employing the instrument developed by Babina et al. (2024), which measures firms' ex ante exposure to AI-

strong universities through pre-existing firm-university hiring networks. Our results remain robust under this approach.

Taken together, our findings show that firm-level AI adoption increases innovation and productivity. The accompanying increase in R&D, combined with the absence of significant changes in physical investment and operating costs, is consistent with an innovation channel through which AI adoption contributes to productivity gains. More broadly, our estimates suggest that the productivity gains associated with firm-level AI adoption can translate into meaningful improvements in aggregate productivity.

2 Related Literature

A growing literature studies the economic implications of AI at the firm and aggregate levels. Prior work documents effects of AI on firm growth and product innovation (Babina et al., 2024), patent examination (Zheng, 2025), disclosure strategies (Cao et al., 2023), firm value (Eisfeldt, Schubert, and Zhang, 2025), industry entry and exit dynamics (Lu, Phillips, and Yang, 2024), and employment (Chen and Wang, 2024). However, existing studies offer mixed predictions about AI's impacts on innovation, and comprehensive empirical evidence remains limited.[3] Using a novel firm-level measure that directly captures the timing of AI adoption, we show that adopters systematically reshape their innovation activities, increasing both the quantity and quality of patenting while altering the novelty, direction, and reach of their innovation.

Our findings also contribute to a broader debate on whether AI fundamentally reshapes the innovation production function with implications for long-run economic and technological

[3]One perspective is that AI can raise innovation by making more information and knowledge usable in R&D through unstructured data processing (Ludwig and Mullainathan, 2024) and automated retrieval and combination of knowledge across technological domains (Agrawal, McHale, and Oettl, 2019), with related benefits from patent text and claim analytics (Setchi et al., 2021; Alderucci and Sicker, 2019). The other perspective is that AI can weaken innovation when complementarities and coordination frictions impede effective deployment (Bresnahan and Trajtenberg, 2017), when control of critical data and algorithms concentrates innovation among early leaders (Cockburn, Henderson, and Stern, 2018), and when AI-enabled prediction reinforce existing technological trajectories (Agrawal, McHale, and Oettl, 2024). AI-generated prior art may flood the system with superficially plausible but poorly disclosed documents that make otherwise patent-worthy inventions unpatentable and reduce firms' incentives to invest in innovation (Ouellette, Fang, and Ouellette, 2025; Yordy, 2021).

trajectories. Some studies argue that advances in AI enable new approaches to scientific and technical research, effectively serving as a general-purpose method of invention that can reshape the innovation process (Cockburn, Henderson, and Stern, 2018). Other work argues that aggregate productivity gains from AI may remain modest if its development primarily focuses on automating existing tasks, and that larger gains require redirecting AI development toward creating new tasks and generating new ideas (Acemoglu and Restrepo, 2020; Acemoglu, 2025). We find that AI adoption is associated with greater efficiency in the use of innovative inputs, as well as changes in inventor composition and more effective knowledge recombination within firms. These patterns are consistent with AI reshaping how firms transform innovative inputs into outputs.

Research on intangible capital shows that modern firms invest heavily in non-physical assets such as R&D, software, and organizational capabilities, and that capitalizing these expenditures makes intangibles a large and growing component of the capital stock and an important driver of growth and productivity (Corrado, Hulten, and Sichel, 2009). A related literature shows that intangible capital constitutes an increasingly important component of firm value and that incorporating intangibles can materially affect measures of risk and valuation (Crouzet et al., 2022; Eisfeldt, Kim, and Papanikolaou, 2022; Kogan and Papanikolaou, 2019). Our findings complement this literature by linking realized AI adoption to subsequent changes in R&D investment, innovation, productivity, and firm value, providing new evidence on the role of AI adoption in firms' intangible investment and value creation.

Finally, and importantly, our paper is also related to recent work investigating the impact of AI on aggregate and firm-level productivity. Existing studies typically extrapolate aggregate effects either by calibrating a task-based framework using the GDP share of tasks impacted by AI and estimated average cost savings in those tasks or a historical analogy to earlier general-purpose technology waves. Under the task-based approach, Aghion and Bunel (2024) report implied increases in annual aggregate TFP growth ranging from 0.07 to 1.24 percentage points (pp), with a median estimate of 0.68 pp per year, while the preferred calibration in

Acemoglu (2025) implies roughly 0.064 pp per year. Under the historical analogy, Aghion and Bunel (2024) benchmark AI against the European electricity wave and the US information and communication technology (ICT) wave, implying about 1.3 and 0.8 pp higher annual productivity growth, respectively, over the next decade. Using U.S. Census manufacturing data, Alderucci et al. (2024) document that firms with at least one AI-related patent have about 8% higher firm-level TFP. We find that AI-adopting firms experience about 8.87% higher value-added TFP on average after adoption than non-adopters, implying a contribution of approximately 0.61% to aggregate value-added TFP growth in a typical post-adoption year, after accounting for offsetting effects on non-adopters.

3 Data and Variables

3.1 Patent Data and Sample

To measure firms' innovation activities, we collect data on patents, citations, technology classifications, and inventors from the United States Patent and Trademark Office (USPTO) using the May 5, 2025 data release. Firm financial data are obtained from CRSP/Compustat Merged. We exclude firms in financial industries (SIC code 6000 to 6999)[4] and firm-year observations with missing revenue data or less than $10 million in assets. Firms are required to have filed at least one patent during the sample period. For these firms, we retain firm-year observations with and without patent filings to mitigate concerns about selection bias.

To link patents to CRSP/Compustat merged firms, we construct a comprehensive linking table by integrating link tables from Kogan et al. (2017), Wharton Research Data Services

[4]We exclude firms in financial industries because their business models and balance sheets differ systematically from those of non-financial firms, making it difficult to account for differences in observable characteristics. Moreover, our estimation of value-added TFP relies on a Cobb–Douglas production function, for which the conditions required to interpret the residual as technological efficiency are less likely to hold in the financial sector due to the nature of financial intermediation. In particular, measured value added reflects not only quantities of services but also but also variation in interest margins, risk premia, and asset prices, which embed pricing and risk-taking components, while balance sheet expansion and leverage can increase output without corresponding changes in productivity. In addition, key inputs such as capital are endogenously determined by regulatory constraints and risk choices, complicating the stable production mapping assumed by the model. Consequently, the residual may conflate productivity with risk-taking and financial structure, thereby limiting its economic interpretability.

(WRDS), and Stoffman, Woeppel, and Yavuz (2022). We further supplement this table using fuzzy name matching and manual verification through SEC filings, Google searches, D&B Hoovers, PR Newswire, Wikipedia, and other sources. To enhance matching accuracy, we use additional firm-specific information, such as locations, prior M&A activities, and name change histories to identify the most likely firm match. We only retain those matches for which we have high confidence.

The final sample consists of 2,600 firms and 791,040 granted patents filed during the period 2011 to 2021. We stop the sample in 2021 to minimize truncation bias caused by the delay between patent application and grant, as patents usually take 2 to 3 years to be granted (Hall, Jaffe, and Trajtenberg, 2001). All variables are winsorized at 1% and 99%. Following the innovation literature, missing values in innovation measures are set to zero to reflect a lack of output rather than missing data (Hall, Jaffe, and Trajtenberg, 2001).

3.2 AI Adoption Data and Measure

We use detailed site-level IT installation records for U.S. firms from the Spiceworks Ziff Davis (SWZD) Historical Company Intelligence Database to measure firm's adoption of AI. This database provides comprehensive information on IT products and their installation dates across firm sites. SWZD compiles data from multiple sources on firms' installations of software, hardware, networks, and telecommunications equipment. It collects information through interviews with senior IT executives and supplements these data with systematic collection methods, including client engagement, web scraping, job postings, document analysis, and purchasing customer lists from vendors. The data are updated annually and undergo extensive quality checks to ensure accuracy. According to SWZD, this database captures approximately 90% of global IT purchasing activity, offering a comprehensive view of technology adoption. The specificity and granularity of this data make it well-suited for tracking the timing of AI adoption across firms. We aggregate site-level technology installation data to the firm level and link them to historical company names from CRSP for Compustat firms using a fuzzy matching

algorithm. All matches are manually reviewed to ensure accuracy. For cases where names do not match exactly, we conduct internet searches and include the observation only when we are confident in the validity of the match.

We construct a firm-level measure of AI adoption by identifying the earliest observed installation date of an AI product across the firm's sites.[5] The first installation year defines the onset of AI adoption. Because the measure is based on observed AI product installations, it captures realized adoption and the introduction of AI capabilities into firm operations. The timing of these installations allows us to trace changes in firm outcomes around AI adoption. A firm is classified as an AI-adopting firm if it adopts any AI technology during the sample period and as a non-adopter otherwise.

3.3 Innovation Measures

We measure innovation outcomes using patent-based innovation metrics. The filing date of a granted patent serves as the primary indicator of innovation activity, as it most closely reflects the timing of the underlying inventive efforts (Lerner and Seru, 2022).

3.3.1 Quantity and Quality

To capture the quantity of innovation outcomes, we measure the number of patents filed by a firm in a given year. We proxy for patent quality using the number of forward citations, a widely used measure of a patent's technological impact (Hall, Jaffe, and Trajtenberg, 2005). Since forward citations accumulate over time and are therefore subject to truncation bias, we follow Hall, Jaffe, and Trajtenberg (2001) and Lerner and Seru (2022) and address this bias by normalizing each patent's forward citation count by the average citation count of patents filed in the same year and

[5] Examples of such products in our data include machine learning frameworks and platforms such as TensorFlow, scikit-learn, and H2O.ai, enterprise and cloud machine-learning platforms such as Microsoft Azure Machine Learning, Google Cloud Machine Learning Engine, Amazon Machine Learning, and the enterprise AI platform Dataiku, as well as more specialized decision-support systems such as Nara Logics and data-labelling and human-in-the-loop platforms such as CrowdFlower. These examples are illustrative rather than exhaustive. Since our sample ends in 2021, chatbot-based generative AI such as ChatGPT is not included. However, the machine learning frameworks and platforms covered in the database can be used to develop generative AI applications.

technology class.

3.3.2 Innovation Characteristics

3.3.2.1 Innovation Novelty: Originality and Distance

Originality captures the diversity of prior knowledge underpinning the invention (Acharya and Xu, 2017). Following Trajtenberg, Henderson, and Jaffe (1997), we calculate originality as one minus the Herfindahl–Hirschman Index (HHI) of the technology classes across a patent's backward citations:

$$Originality_p = 1 - \sum_{k}^{n_p} B_{pk}^2 \tag{1}$$

where B_{pk} is the number of backward citations made by patent p to patents in technology class k, divided by the total number of backward citations made by patent p. n_p is the number of unique technology classes cited by patent p. A higher originality score indicates that a patent draws knowledge from a more diverse set of technological fields.

Distance measures the technology propinquity between a patent and the overall technology expertise of its owning firm (Akcigit, Celik, and Greenwood, 2016). To construct this measure, we first compute the distance between technology classes, defined by the first two digits of the Cooperative Patent Classification (CPC) code. Specifically, the distance metric between technology classes X and Y is calculated based on the citations made to classes X and Y:

$$d(X,Y) \equiv 1 - \frac{\#(X \cap Y)}{\#(X \cup Y)} \tag{2}$$

where $\#(X \cap Y)$ denotes the number of patents that cite at least one patent from both class X and Y (intersection) and $\#(X \cup Y)$ denotes the number of patents that make citations to at least one patent from either class X or Y (union). If two classes are frequently cited together, they are considered technologically closer. Conversely, if they are rarely co-cited, the classes are deemed

more distant.[6]

Next, we calculate the distance of a patent p to its filing firm i, proxied by the average distance between p and all other patents previously filed by firm i. The metric is computed as follows:

$$d_t^{\iota}(p,i) \equiv \left[\frac{1}{\|\mathcal{P}_{it}\|} \sum_{p' \in \mathcal{P}_{it}} d(X_p, Y_{p'})^{\iota} \right]^{\frac{1}{\iota}} \tag{3}$$

where $\mathcal{P}_{it}$ denotes the set of all patents filed by firm i prior to the patent p in year t. The cardinality $\|\mathcal{P}_{it}\|$ represents the size of firm i's patent portfolio accumulated before filing patent p in year t. $d(X_p, Y_{p'})$ denotes the distance between a patent p in technology class X and another patent p' in technology class Y, where p' represents each other patent in firm i's portfolio. Following Akcigit, Celik, and Greenwood (2016) and Ma, Tong, and Wang (2022), we set ι, the power parameter of the generalized mean operator, equal to 2/3. Here, $d_t^{\iota}(p,i)$ denotes the technology distance between patent p to its owning firm i in year t. A higher value indicates that the patent lies in a more technologically distant area relative to the firm's existing technological expertise.

3.3.2.2 Innovation Direction: Exploitative and Explorative

Exploitative patents refer to inventions that mainly build on a firm's existing knowledge base. Following Benner and Tushman (2002) and Gao, Hsu, and Li (2018), we define a firm's existing knowledge base for a given patent as (i) all patents filed by the firm and (ii) all citations made by the firm's patents from the beginning of the fifth calendar year prior to that patent's filing date to its filing date. We classify a patent as exploitative if at least 80% of its citations draw on this existing knowledge base (that is, either belong to the firm's own patent filings or were cited by the firm during this period). This measure captures the extent to which a firm's innovation improves along the existing technology trajectory (Benner and Tushman, 2002).

[6]For robustness, we compute two versions of the distance metric for a given class pair in a given year: one based on citations in that year alone and another based on cumulative citations up to that year. The two approaches yield similar results. We report results using technology distances calculated from citations accumulated up to the given year.

Explorative patents refer to inventions that are mainly built outside of a firm's existing knowledge base. A patent is classified as explorative if at least 80% of its citations are outside the firm's existing knowledge base.[7] This measure gauges the extent to which a firm's innovation shifts to a new technology trajectory (Benner and Tushman, 2002).

3.3.2.3 Innovation Reach: Generality and Patent Claim

Generality reflects how widely a patent's impact spreads across different technology classes (Hall, Jaffe, and Trajtenberg, 2001). Following Trajtenberg, Henderson, and Jaffe (1997), we calculate generality as one minus the HHI of the technologies classes in the forward citations citing a patent:

$$Generality_p = 1 - \sum_{k}^{n_p} F_{pk}^2 \tag{4}$$

where F_{pk} represents the ratio of forward citations made to patent p from patents in technology class k to the number of total forward citations made to patent p across all technology classes. n_p is the number of unique technology classes citing patent p. A higher generality score indicates that a patent influences subsequent patents in a wider variety of technology domains.

Patent claims define the scope of a patentee's exclusive rights and are of significant economic and legal importance (Lerner, 1994; Lee, 2010). Following Akcigit, Baslandze, and Stantcheva (2016) and Aghion et al. (2019), we use the number of claims in a patent application to measure the scope of legal protection.

4 Empirical Analysis

4.1 Summary Statistics

Table 1 reports innovation outcomes and firm characteristics for the full sample, AI adopters, and never-adopting firms. Out of 2,600 firms in our sample, 518 (20%) firms adopt AI technologies,

[7] As a robustness check, we also use a 60% threshold to define exploitative and explorative patents and find similar results.

while 2,082 (80%) firms never do so. Appendix Table B.2 reports the industry distribution of AI-adopting firms using the Fama-French 49 classification. The largest numbers of adopters are in computer software, electronic equipment, retail, business services, and pharmaceuticals, while the highest adoption rates are in transportation, entertainment, retail, utilities, and measuring and control equipment. Appendix Figure B.1 shows the geographic diffusion of AI adoption across U.S. BEA Economic Areas from 2016 to 2021.[8] Adoption is initially concentrated in areas such as San Francisco, New York, Washington D.C., Chicago, and Los Angeles. Over time, these hubs deepen their adoption and additional areas, including several in Texas and the Midwest, begin to exhibit noticeable numbers of adopters. By 2021, high adoption remains clustered in major coastal and metropolitan areas, but medium levels of adoption are visible across much of the country, indicating gradual yet uneven spatial diffusion of AI.

Table 1 shows that AI-adopting firms produce more patents, and those patents, on average, receive more adjusted citations, exhibit greater originality and technology distance, are more exploitative, and contain more claims. AI-adopting firms exhibit higher labor, but lower value-added TFP on average and capital compared with never-adopting firms.[9] It is worth noting that AI-using firms, on average, spend less on R&D than never-using firms, and the difference is both economically and statistically significant. This pattern alleviates the concern that our results are driven by firms that systematically invest more in R&D. We find that AI-adopting firms are larger, have higher *Tobin's q*, and hold less cash compared to never-adopting firms.

4.2 Baseline Model

Since firms adopt AI at different times, traditional two-way fixed effects models can suffer from the "bad comparisons" problem when treatment effects vary across cohorts or over time (Baker,

[8]The earliest adoption year in our sample is 2016, consistent with evidence that AI adoption accelerated after 2015 (Acemoglu et al., 2022). This timing also coincides with major advances in widely used deep-learning tools and architectures, including TensorFlow's release in November 2015 and the Transformer architecture introduced in 2017.

[9]Value-added TFP is constructed as the residual from a log Cobb–Douglas value-added production function. Negative values indicate that, conditional on inputs, firms produce less value added than predicted by the estimated production function. Construction details are provided in Appendix Table B.1.

Larcker, and Wang, 2022). To address this concern, we follow Gormley and Matsa (2011, 2016) and employ a stacked difference-in-differences design to estimate the effects of AI adoption on firm-level outcomes. This approach compares treated and control firms within each cohort while controlling for time-invariant firm heterogeneity and shocks common to firms in the same cohort and year. We define each cohort by its AI adoption year, which we refer to as the event year. For each cohort, we construct an event window spanning five years before to five years after the event year. Treated firms are those that adopt AI in the event year, while the control group includes: (1) firms that never adopt AI during the sample period (never-treated), and (2) pre-adoption observations of firms that adopt AI later (not-yet-treated). We then stack all cohorts of treated and control firms to form the stacked sample.

To examine how AI adoption affects innovation, we estimate the following model using the stacked sample separately for each horizon $h \in \{1, 2, 3\}$:

$$\mathbb{E}\left[\, Y_{i,j,c,t+h} \mid \cdot \,\right] = \exp\Big(\beta_h \, AI_{i,j,c,t} \;+\; \alpha_{i,c} \;+\; \psi_{t+h,c}\Big) \tag{5}$$

where i, j, c, t denote firm, industry, cohort, and year, respectively. $Y_{i,j,c,t+h}$ is the innovation measure for firm i in year $t+h$. $AI_{i,j,c,t}$ equals 1 if firm i in cohort c adopts AI in year t and 0 otherwise. The estimations include year-cohort fixed effects ($\psi_{t+h,c}$), which absorb calendar-year shocks common to firms within each cohort, including aggregate economic conditions and time trends, and firm-cohort fixed effects ($\alpha_{i,c}$), which absorb time-invariant firm characteristics within each cohort, such as persistent differences in management quality, corporate culture, and technology. Industries are defined based on the Fama–French 49 classification. Standard errors are clustered at the firm level. Since the innovation outcomes are non-negative with many zeros, we estimate Equation (5) by Poisson pseudo-maximum likelihood (PPML), following Cohn, Liu, and Wardlaw (2022). The coefficients of interest, β_h, capture the differential changes in innovation outcomes between AI-adopting firms and control firms within the same cohort at each event-time horizon h.

4.3 Baseline Results

4.3.1 Innovation Quantity and Quality

We first examine the effect of AI adoption on innovation quantity, measured by the number of patents. Panel A of Table 2 reports positive and significant coefficients on AI, indicating that AI adopters increase post-adoption patenting relative to non-adopters. The coefficients are also economically significant. Three years after adoption, AI-adopting firms generate approximately 19.1% ($(e^{0.175} - 1) \times 100\%$) more patents than non-adopting firms.

Next, we examine how AI adoption influences the quality of innovation. Specifically, we evaluate the scientific value of patents, measured by forward citations adjusted for technology class and application year. As shown in Panel B of Table 2, the coefficients on AI are all positive and significant, indicating that patents produced by AI-adopting firms receive more adjusted citations after AI adoption than those from non-adopting firms. The result suggests that AI adopters develop patents of higher scientific value after using AI than non-adopters.

4.3.2 Innovation Characteristics

We further study whether the use of AI affects the innovation characteristics of AI-adopting firms. Table 3 Panel A reports the results for the impacts of AI on innovation novelty. We observe that AI-using firms experience a significant increase in originality and technology distance of their patents following adoption, compared to non-adopting firms. The results indicate that AI-adopting firms subsequently produce patents that draw on diverse set of technological fields and are more distant from their prior patent portfolios.

Panel B of Table 3 focuses on how AI adoption affects innovation direction. The coefficients on AI are all positive and significant, indicating that, following AI adoption, adopting firms generate more exploitative and explorative patents than non-adopting firms. These results suggest that AI adoption is associated with greater exploitation of existing technological knowledge as well as greater exploration of new technological knowledge.

We also assess whether AI adoption broadens the reach of firms' innovation. Columns (1)-(3) of Table 3 Panel C show that following AI adoption, patents produced by AI-adopting firms exhibit significantly higher generality than those of non-adopting firms, indicating that these patents are relevant to a more diverse range of subsequent innovations. Columns (4)-(6) show that AI-adopting firms also produce patents with significantly more claims, compared to non-using firms. Three years after adoption, their patents contain approximately 17.48% more claims than those of non-using firms. These results indicate that AI adoption expands both the technological reach and legal scope of firms' innovation.

4.4 Dynamic Effects

The baseline estimates indicate that AI adopters subsequently outperform non-adopters in innovation. These estimates may be biased by reverse causality, whereby firms experiencing improvements in innovation are more likely to adopt AI, or by differential pre-trends if adopters and non-adopters follow distinct pre-adoption trajectories for reasons unrelated to AI. To address these concerns, we augment the stacked DiD model with a cohort-specific event study. For each adoption cohort, we estimate event-time coefficients (leads and lags relative to the year before adoption) to test pre-adoption parallel trends and to trace the timing and persistence of post-adoption effects on innovation.

Following Baker, Larcker, and Wang (2022), let $t_{i,c}$ denote firm i's adoption year in cohort c and define event time as $\ell \equiv t - t_{i,c}$. For each ℓ, the indicator $AI^{\ell}_{i,j,c,t} = \mathbf{1}\{t - t_{i,c} = \ell\}$, which equals 1 if t is ℓ years relative to adoption and 0 otherwise. We estimate dynamic effects using a cohort-specific event-study:

$$\mathbb{E}\big[\, Y_{i,j,c,t+1} \mid \cdot \,\big] = \exp\left(\sum_{\ell=-5}^{-2} \beta_\ell \, AI^{\ell}_{i,j,c,t} \;+\; \sum_{\ell=0}^{5} \beta_\ell \, AI^{\ell}_{i,j,c,t} \;+\; \alpha_{i,c} \;+\; \psi_{t+1,c} \right) \tag{6}$$

where, $Y_{i,j,c,t+1}$ is the innovation outcome for firm i in year t+1. We use $\ell = -1$ (one year before adoption) as the reference period. $\alpha_{i,c}$ and $\psi_{t+1,c}$ are firm-by-cohort and year-by-cohort fixed

effects. Standard errors are clustered at the firm level. This model uses a set of relative-time indicators to capture the time periods before the treatment ("leads") and the time periods after treatment ("lags").

Figure 1 reports the coefficients from the lead-lag regressions. Panel A shows that the difference in patent quantity and quality between AI-adopting and non-adopting firms increases significantly three years after AI adoption, accounting for the one-year-ahead measurement of the outcomes. Panels B, C, and D examine patent novelty, direction, and reach, respectively. Differences between AI adopters and non-adopters emerge after adoption for originality, distance, exploitative patents, generality, and claims, while explorative patents show no significant changes. We find no significant pre-adoption differences, mitigating concerns about reverse causality or differential pre-trends.

Our results show that the impact of AI adoption on corporate innovation emerges relatively quickly. This is consistent with our use of patent filing rather than grant dates to measure innovation output. Moreover, industry evidence documents that AI-enabled R&D projects often progress on the scale of months rather than years (Microsoft, 2024; Chu, 2025). Prior work also indicates that the transition from R&D to patenting is relatively rapid and that firms tend to apply for patents early in the R&D process rather than upon its completion (Hall, Griliches, and Hausman (1984); Griliches (1998); Moretti (2021); He and Qiu (2025)).

4.5 Addressing Endogeneity

One concern with the baseline results is endogeneity arising from non-random selection into AI adoption. The inclusion of firm–cohort fixed effects mitigates selection driven by time-invariant differences across firms within each cohort stack. However, endogeneity may persist if time-varying, unobserved changes, such as improvements in management quality or shifts toward a more innovation-oriented strategy, simultaneously increase firms' propensity to adopt AI and their innovation outcomes, even in the absence of AI adoption.

To alleviate this concern, we implement a control function regression approach using an

instrumental variable to isolate exogenous variation in AI adoption. In the first stage, AI adoption is estimated as a function of the instrument and other exogenous covariates, and generalized residuals are obtained from the estimation. In the second stage, these residuals are included as an additional control to account for endogeneity. Compared with two-stage least squares (2SLS), the control function approach is better suited to nonlinear models with endogenous explanatory variables (Wooldridge, 2015), making it appropriate for our setting.

Following Babina et al. (2024), we use *AI exposure* as an instrumental variable for AI adoption. The instrument captures a firm's *ex ante* exposure to AI talent from historically AI-strong universities through its pre-existing STEM hiring network as of 2010. The relevance condition follows from the importance of specialized human capital for AI adoption. Firms with greater access to AI-skilled workers face fewer constraints in adopting AI and should therefore be more likely to adopt AI (Babina et al., 2024).

Several considerations support the plausibility of the exclusion restriction. *AI exposure* is constructed from firms' predetermined hiring networks and universities' AI research strength prior to 2010, mitigating concerns that it reflects anticipation of future AI adoption. To address the possibility that exposure to AI-strong universities captures access to high-quality human capital more generally, we also control for *Top 10 exposure*, which measures firms' exposure to top 10 universities through the same 2010 STEM hiring networks (Babina et al., 2024).[10]

To implement the control-function approach, we first estimate the following probit model on the stacked panel:

$$\Pr(AI_{i,j,c,t} = 1 \mid z_{i,j,c,t}) = \Phi\left(\beta\, AI\ exposure_{i,j,c} + \theta\, Top\ 10\ exposure_{i,j,c} + \gamma_{j,c} + \tau_{t,c}\right), \quad (7)$$

where $AI\ exposure_{i,j,c}$ measures exposure to AI-strong universities through firm-university STEM hiring networks as of 2010 (Babina et al., 2024), and $Top\ 10\ exposure_{i,j,c}$ measures exposure to top-10 universities through the same networks. $\gamma_{j,c}$ and $\tau_{t,c}$ denote industry-by-

[10] AI strength does not simply reflect overall university quality: among the 44 AI-strong universities, 11 are ranked among the top 20 and 25 among the top 50 by U.S. News & World Report. Appendix Table B.3 provides the complete lists of the top 50 universities and AI-strong universities identified by Babina et al. (2024).

cohort and year-by-cohort fixed effects, respectively. We exclude firm fixed effects because they would absorb the cross-sectionally varying instrument.

Let $z_{i,j,c,t}$ denote the vector of first-stage regressors and $\hat{\boldsymbol{\delta}}$ the corresponding coefficient estimates. We construct the generalized residual as

$$\widehat{Residual}_{i,j,c,t} = AI_{i,j,c,t}\frac{\phi\left(z'_{i,j,c,t}\hat{\boldsymbol{\delta}}\right)}{\Phi\left(z'_{i,j,c,t}\hat{\boldsymbol{\delta}}\right)} - \left(1 - AI_{i,j,c,t}\right)\frac{\phi\left(z'_{i,j,c,t}\hat{\boldsymbol{\delta}}\right)}{1 - \Phi\left(z'_{i,j,c,t}\hat{\boldsymbol{\delta}}\right)}, \tag{8}$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ denote the standard normal PDF and CDF, respectively.

In the second stage, we estimate the PPML model including the generalized residuals:

$$\mathbb{E}[Y_{i,j,c,t+h} \mid \cdot] = \exp\left(\beta\, AI_{i,j,c,t} + \theta\, Top10\, exposure_{i,j,c} + \kappa\, \widehat{Residual}_{i,j,c,t} + \gamma_{j,c} + \psi_{t+h,c}\right) \tag{9}$$

where $Y_{i,j,c,t+h}$ is the innovation outcome for firm i in year $t + h$. $\gamma_{j,c}$ and $\psi_{t+h,c}$ are industry-by-cohort and year-by-cohort fixed effects, respectively. Standard errors are computed using a firm-level cluster bootstrap with 1,000 replications.

Table 4 reports the control function regression results. For brevity, we report the estimates in year $t + 3$, as the results for other years are similar.[11] Column (1) presents the results of the first-stage regression. The estimated coefficient on *AI exposure* is positive and significant at the 1% level, indicating that exposure to AI-strong universities is positively associated with firms' likelihood of adopting AI technology. The significant Wald χ^2 statistic ($\chi^2 = 541.82$, $p < 0.001$) indicates that the instrument strongly predicts the endogenous variable in the first stage, supporting the relevance condition. Columns (2)-(9) present the second-stage regression results. The positive effects of AI adoption persist across various innovation outcomes, including *Patent count*, *Adjusted citations*, *Originality*, *Distance*, *Exploitative*, *Generality*, and *Claim*. The negative coefficients on the generalized residuals indicate negative selection, i.e., factors

[11] The probit model in the first stage is estimated by maximum likelihood, with convergence attained at the maximum of the log-likelihood function. Estimation may fail when the binary outcome is perfectly or nearly perfectly predicted by the covariates, or when the within-firm sample size is insufficient for reliable identification. Consequently, firms with very few observations are excluded from the estimation sample, which reduces the number of observations available for inference.

increasing the propensity to adopt AI are associated with lower innovation outcomes, biasing baseline estimates downward. Accounting for this selection leads to larger estimates relative to the baseline. Overall, these results suggest that differences in innovation quantity, quality, and characteristics between AI-adopting and non-adopting firms remain robust after accounting for the endogeneity of AI adoption.

4.6 Robustness Checks

We also conduct several robustness checks of our baseline results. First, we use entropy balancing to address observable differences between AI-adopting and non-adopting firms following Hainmueller (2012). Second, we conduct a placebo test in which fictitious adoption years are randomly assigned to firms that never adopt AI. Third, we restrict the control group to never-treated firms to address potential anticipatory behavior among future adopters. Finally, we exclude AI-related patents identified by Giczy, Pairolero, and Toole (2022) to examine whether the documented effects extend beyond AI-specific inventions. As reported in Appendix Table B.4, the entropy-balancing, alternative control group, and non-AI patent specifications yield results consistent with our baseline findings, while the placebo tests produce no corresponding positive effect. Additional details on the implementation of these tests are provided in the Appendix A.

5 Channels

The results thus far suggest that AI adoption significantly affects firms' innovation. We next examine the potential channels underlying these effects.

5.1 Efficiency Channel

5.1.1 Inventor Productivity

We begin by examining an efficiency channel through which AI can increase innovation output from a given set of inventive inputs. We first focus on inventor productivity. AI may increase patent output per inventor by automating time-intensive tasks, improving information processing, and facilitating the diffusion of knowledge and effective practices within firms (Setchi et al., 2021; Brynjolfsson, Li, and Raymond, 2025). To test this channel, we define *Inventor productivity* as the average number of patents filed per inventor at firm i in year t. We first calculate each inventor's annual patent output, assigning each inventor an equal fractional share of jointly invented patents following Moretti (2021), and then average inventor productivity within each firm-year.[12] Columns (1)-(3) of Panel A in Table 5 report estimates from the baseline specification in Equation (5). Inventor productivity increases significantly following AI adoption relative to non-adopting firms, reaching 17.59% higher productivity three years after adoption. These results indicate that AI promotes innovation through increased inventor productivity.

5.1.2 Innovation Efficiency

We next examine innovation efficiency, which captures how effectively firms convert R&D spending into patents. AI may improve innovation efficiency by enhancing project selection and reducing the cost of experimentation. AI-based prediction can improve the evaluation of research opportunities and the allocation of resources toward more promising projects (Cockburn, Henderson, and Stern, 2018; Yoo, Jung, and Jun, 2023). AI can also accelerate experimentation by improving candidate screening and enabling exploration of larger search spaces, as demonstrated in applications to materials and drug discovery (Merchant et al., 2023; Bennett et al., 2025). These capabilities can enable firms to generate more innovation output

[12] As a robustness check, we measure *Inventor productivity* as the number of patents filed by firm i in year t divided by the number of its inventors who file patents that year. The results are similar.

from a given level of R&D spending.

To test this channel, we define *Innovation efficiency* as the ratio of the number of patents to R&D expenditure following Gao and Chou (2015). We estimate the baseline model (Equation (5)) using *Innovation efficiency* as the outcome. Columns (4)-(6) of Panel A in Table 5 show that innovation efficiency increases significantly following AI adoption relative to non-adopting firms, with an estimated increase of 18.07% three years after adoption. These results indicate that AI promotes innovation by enabling firms to generate more patents from a given level of R&D spending.

5.2 Inventor Composition Channel

AI adoption may affect not only inventor productivity but also the size and composition of inventor teams. More broadly, the technology and labor literature shows that the adoption of information technology can be accompanied by changes in workplace organization and labor demand (Bresnahan, Brynjolfsson, and Hitt, 2002). AI may be particularly relevant for inventive labor because it can reshape the innovation process and the organization of R&D (Cockburn, Henderson, and Stern, 2018).

The effect of AI adoption on inventor team size is theoretically ambiguous. If AI substitutes for tasks previously performed by inventors, firms may require fewer inventors for a given level of innovation. In contrast, if AI complements inventive labor by raising inventor productivity or expanding the scope of R&D, firms may increase their innovation efforts and employ more inventors. AI adoption may also lower barriers to participation in inventive activity. By making some components of the innovation process easier to perform, AI may enable employees without prior inventive experience to contribute to invention, thereby increasing the entry of first-time inventors. We therefore examine whether AI adoption is associated with changes in inventor team size and the participation of first-time inventors.

To explore this channel, we define *Team size* as the firm-year average number of inventors per patent. *First-time inventor* is defined as the average share, across the firm's patents filed in year t,

of inventors in each patent's inventor team who have no prior USPTO patent filings before year t. Using these measures as dependent variables, we re-estimate the stacked DiD model to compare changes in outcomes between AI-adopting and non-adopting firms following AI adoption. Panel B of Table 5 shows that AI-adopting firms experience significant increases in *Team size* and *First-time inventor* relative to non-adopting firms. These results indicate that, following AI adoption, AI-adopters form larger inventor teams and have a higher share of inventors with no prior patenting experience than non-adopters.

5.3 Knowledge Recombination Channel

One view in the innovation literature conceptualizes invention as a recombinant process in which existing ideas and technological components are combined in novel ways (Weitzman, 1998; Fleming, 2001). Recombinant invention depends on the ability to identify relevant prior knowledge, search over potential combinations, and evaluate their viability under technological uncertainty (Fleming, 2001). As the stock of knowledge expands, however, it becomes increasingly difficult for individual inventors to span the knowledge frontier, leading inventors to specialize and rely more heavily on teams that combine complementary expertise (Jones, 2009; Wuchty, Jones, and Uzzi, 2007).

We propose that AI can facilitate knowledge recombination by reducing the costs of searching, evaluating, and integrating technological knowledge. By making it easier to identify relevant knowledge and assess potential combinations, AI can facilitate the integration of specialized expertise across technological domains. This mechanism can affect how knowledge is organized within inventor teams. Individual inventors can specialize more deeply in their areas of expertise while teams combine inventors with core skills spanning a broader set of technological domains. Such an organizational structure preserves depth of expertise at the inventor level while expanding the breadth of knowledge available for recombination at the team level.

To examine this mechanism, we consider two complementary dimensions of inventive

knowledge. *Specialization* captures the concentration of individual inventors' prior patenting experience across technological fields, with higher values indicating greater specialization. *Core skill breadth* captures the number of distinct core technological skills represented within an inventor team. Using these measures as dependent variables, we re-estimate Equation (5) to examine how knowledge recombination changes following AI adoption relative to non-adopting firms. Panel C of Table 5 shows that, relative to non-adopting firms, AI-adopting firms experience significant increases both *Specialization* and *Core skill breadth*. Following AI adoption, firms increasingly draw on inventors with more concentrated technological expertise while assembling teams with a broader range of core technological skills. This combination of individual specialization and team-level breadth is consistent with AI facilitating the recombination of specialized knowledge. Together, these findings provide evidence that AI reshapes the process of innovation itself.

6 Cross-Sectional Heterogeneity

6.1 Business Focus

AI may be particularly valuable in complex information environments, where firms must process and integrate information across a wider range of activities. By improving prediction and information processing, AI can generate greater gains when decisions involve many potential states (Agrawal, Gans, and Goldfarb, 2018). We therefore conjecture that the effects of AI adoption are stronger for diversified firms, which operate across more heterogeneous activities and face greater information-processing demands.

To test this hypothesis, we examine whether the effect of AI on innovation differs between focused and diversified firms. Following Comment and Jarrell (1995), we use the number of business segments to measure a firm's business focus at the firm-year level. Specifically, we construct an indicator variable called *Focused firm*, which equals 1 if firm i has only one business segment in 2015, and 0 otherwise. The measure is constructed using 2015 values, one

year before the earliest adoption date in the sample, to preclude contamination from changes in business focus induced by AI. We include an interaction term between *AI* and *Focused firm* in Equation (5). To alleviate the concern that time-invariant firm characteristics and/or time trends may differ systematically between firms with different business focus, we include firm-cohort-focused firm and year-cohort-focused firm fixed effects.

Table 6 Panel A reports the results on innovation quantity, quality, and characteristics in the third year following AI adoption. The positive and statistically significant coefficients on *Originality*, *Distance*, *Exploitative*, and *Claim* indicate that AI adoption has a particularly pronounced effect on these aspects of innovation among firms with a focused business scope. These results are inconsistent with the hypothesis that AI disproportionately benefits diversified firms with more complex information environments by enabling them to process and exploit large volumes of information.

Several mechanisms may explain why focused firms derive greater innovation benefits from AI. First, a narrower business scope may generate more coherent internal data, improving AI performance (Zha et al., 2025; Whang et al., 2023; Brynjolfsson and McAfee, 2017). Second, focused firms may possess deeper domain expertise, an important complement to AI (Brynjolfsson, Rock, and Syverson, 2019; Azoulay, Krieger, and Nagaraj, 2025). Third, focused firms may face fewer resource-allocation and coordination frictions when making complementary investments and integrating AI into inventive activities (Rajan, Servaes, and Zingales, 2000). These mechanisms offer potential explanations for the stronger effects of AI adoption among focused firms.

6.2 Firm Size

A natural concern is that the estimated effect of AI adoption on innovation is confounded by firm size. Larger firms may systematically differ from smaller ones, both in their propensity to adopt AI and in the benefits they can extract from it. In particular, large firms have greater financial resources and richer data assets, which can make them more willing and able to adopt

AI technologies, and prior work shows that such firms benefit more from AI investment (Babina et al., 2024). Large firms may inherently generate more innovation regardless of AI adoption, raising the concern that our baseline estimates partly reflect differences associated with firm size rather than AI adoption.

To examine whether our results are driven by firm size, we augment Equation (5) with an interaction term between *AI* and an indicator for *Big firm*, which equals 1 if the total assets of a firm in 2015 are at least as large as the industry average value in 2015, and 0 otherwise. We also include firm-cohort-big firm and year-cohort-big firm fixed effects to account for time-invariant firm characteristics and time-varying trends that differ across firms of different sizes. Panel B of Table 6 shows that the interaction coefficients are statistically insignificant across innovation outcomes, with only weak significance for *Adjusted citation*. Thus, the effects of AI adoption on innovation do not appear to be driven by large firms.

6.3 Financial Constraint

Prior research shows that financial constraints shape firms' investment decisions (Almeida and Campello, 2007; Li, 2011). AI requires large complementary investments in software, data infrastructure, organizational change, and human capital, which entail substantial up-front and recurring costs (Brynjolfsson, Rock, and Syverson, 2021). Financial constraints can prevent firms from making the substantial investments needed for effective AI adoption. Such constraints may also depress innovation by limiting R&D spending and hindering talent acquisition. As a result, observed differences between AI adopters and non-adopters may reflect both technology choices and underlying financial capacity.

To examine whether financial constraints explain weaker innovation among non-adopters compared with AI adopters, we introduce an interaction term between *AI* and *Unconstrained firm* in Equation (5). Using the financial constraint measure developed by Hoberg and Maksimovic (2015), we define *Unconstrained firm* as an indicator variable equal to 1 if a firm's risk of delaying investments due to liquidity issues is not greater than the 2015 industry average, and

0 otherwise. We control for firm-cohort-unconstrained firm and year-cohort-unconstrained firm fixed effects to absorb differences related to firms' financing constraints across cohorts and over time. Table 6 Panel C show that the coefficients on *AI*×*Unconstrained firm* are insignificant across all innovation measures. These results are inconsistent with financial constraints driving the differential effects on innovation between AI adopters and non-adopters.

6.4 Technological Capabilities

Firms differ in their technological capabilities. High-tech firms with more developed R&D infrastructure and inventor expertise, might be more likely to adopt AI and inherently more innovative. Webb et al. (2018) document a rapid increase in patenting among U.S. technology firms in recent years, especially in software, cloud computing, and artificial intelligence. This raises the concern that our results could reflect preexisting differences in technological capabilities rather than the effect of AI adoption.

To alleviate this concern, we augment the baseline model with an interaction term between *AI* and *High-tech firm* to examine whether the effects of AI adoption on innovation differ between high-tech firms and other firms. *High-tech firm* is an indicator variable that equals 1 if firm *i* operates in high-tech sectors in 2015, and 0 otherwise. We classify high-tech sectors as 2-digit NAICS sectors 51 ("Information") and 54 ("Professional, Scientific, and Technical Services") following Babina et al. (2024). Firm-cohort-high-tech firm and year-cohort-high-tech firm fixed effects are included to mitigate concerns that firms with different technological capabilities may systematically differ in their time-invariant characteristics and/or exhibit distinct time-varying trends. Panel D of Table 6 reports that the coefficients on *AI*×*High-tech firm* are mostly insignificant, indicating that high-tech firms do not experience a stronger effect of AI adoption on innovation compared to firms in non-high-tech industries. These results ease the concern that our results are driven by pre-existing differences in technological capabilities between AI-adopting and non-adopting firms.

To further assess whether the effects of AI on innovation are driven by technology-intensive

industries, we exclude firms operating in tech sectors in 2015 from our sample and re-estimate Equation (5). Appendix Table B.5 shows that the impacts of AI on innovation remain significant after excluding high-tech firms. Our results are consistent with Babina et al. (2024), which document that the benefits of AI adoption extend across a broad range of economic sectors, not only those traditionally classified as high-tech.[13]

7 Productivity and Economic Outcomes

The evidence so far indicates that AI adoption promotes innovation by raising efficiency, altering the composition of inventors, and enabling richer knowledge recombination. We now investigate whether these innovation effects are accompanied by changes in firms' investment, operating costs, firm value, and productivity. This analysis helps distinguish an innovation-oriented mechanism involving intangible investment from alternative mechanisms operating through physical capital deepening or immediate cost reductions.

7.1 Intangible versus Tangible Investment

We first examine whether AI adoption is associated with greater intangible investment, measured by R&D, or with higher physical capital investment, measured by capital expenditures. If AI operates mainly by enhancing firms' inventive capabilities and their stock of knowledge capital, we expect to observe stronger adjustments along the R&D margin than in physical investment. If instead AI operates primarily as a physical capital–embodied technology that induces higher physical capital investment, we should observe a response in capital expenditures.

[13] We also examine whether product market competition shapes the impact of AI on innovation. Appendix Table B.6 shows that competitive pressure amplifies the effect of AI adoption on *Originality* and *Exploitative*, but does not significantly affect other innovation outcomes. These results suggest that under intense competition, firms use AI to generate more novel inventions and to strengthen existing technological capabilities. In highly competitive markets, firms face pressure to differentiate and respond quickly (Arts, Cassiman, and Hou, 2025). Such pressure can shift innovation toward exploitative activities that rely on established capabilities, entail lower uncertainty, and yield more immediate returns. Jansen, Van Den Bosch, and Volberda (2006) show that exploitative innovation is more closely associated with short-term financial performance in competitive environments, making it a strategic priority when market pressure is high.

To examine how AI adoption affects the investment of AI-adopters relative to non-adopters, we re-estimate the stacked DiD specification in Equation (5) with *R&D* and *Capital expenditures* as alternative dependent variables. Since these outcome variables are continuous measures, we estimate the specification using a linear model. Table 7 Panel A shows that AI adopters increase investment in R&D following adoption, whereas their capital expenditures do not rise relative to non-adopters. This pattern is consistent with AI adoption being accompanied by greater investment in innovation-related intangible capital rather than physical capital. Together with the patent evidence and the mechanisms documented above, these results suggest that AI adoption not only increases resources devoted to innovation but also enhances the effectiveness of the innovation process.

7.2 Operating Costs

We next examine whether AI adoption operates as a cost-cutting technology, as might occur if firms deploy AI to automate production or reduce operating costs. If AI operates primarily through this channel, operating costs should decline following adoption. By contrast, an innovation-oriented mechanism would be more consistent with increased R&D investment without necessarily generating immediate cost savings.

To examine these predictions, we proxy operating costs using three variables: *Production cost*, *Non-production cost*, and *Total operating cost*. *Production cost* is defined as cost of goods sold scaled by sales. *Non-production cost* is based on main SG&A, defined as selling, general and administrative expenses minus R&D expenses and advertising expenses, scaled by sales.[14] *Total operating cost* is measured as total operating expenses scaled by total sales.[15] Using these measures as dependent variables, we re-estimate Equation (5) to examine how AI adoption affects the operating costs of adopting firms relative to non-adopting firms. Panel B Table

[14] According to Compustat, selling, general and administrative expenses (xsga) include R&D expenses (xrd) and advertising expenses (xad). Following Banker et al. (2019) and Ang et al. (2022), we subtract xad and xrd from xsga to measure non-production, non-R&D, non-advertising operating costs.

[15] As a robustness check, we follow Enache and Srivastava (2018) and scale production costs (cogs), non-production costs (xsga-xrd-xad), and total operating costs (xopr) by total assets. The results are similar.

7 shows that the estimated coefficients on AI are all statistically insignificant, indicating that AI-adopting firms do not experience significant changes in operating costs following adoption relative to non-adopting firms. These results do not support immediate operating-cost reductions as a primary source of gains from AI adoption.

7.3 Firm Value

The innovation-related effects of AI adoption may have implications for firm value. If investors expect AI adoption to enhance firms' innovative capabilities and future growth opportunities, AI adoption should be associated with higher firm valuations. To examine the valuation effects of AI adoption, we use *Tobin's q* and *Total q* as proxies for firm value. *Tobin's q* is defined as the market value of assets divided by the book value of assets. *Total q*, developed by Peters and Taylor (2017), is the firm's market value divided by the sum of its physical and intangible capital stocks. By explicitly incorporating an estimate of intangible capital, *Total q* refines the standard *Tobin's q* measure and better captures firms' investment opportunities, particularly for innovative firms.

Using the same stacked DiD framework, we re-estimate Equation (5) with *Tobin's q*, and *Total q* as alternative dependent variables to compare changes in firm value for AI adopters before and after adoption relative to non-adopters. Panel C in Table 7 shows that the estimated coefficients on AI are positive and significant. Firms adopting AI exhibit significant increases in both *Tobin's q* and *Total q* relative to non-adopting firms. These valuation effects are consistent with investors capitalizing expected future cash flows and growth opportunities associated with AI-enabled innovation.

7.4 Firm-Level Productivity

Furthermore, we examine whether AI adoption is associated with gains in firm-level productivity. By facilitating the development of new products and processes, AI may contribute to productivity growth through innovation (Cockburn, Henderson, and Stern, 2018). Greater

investment in R&D may further support these gains through the accumulation of knowledge capital (Corrado, Hulten, and Sichel, 2009; Peters and Taylor, 2017).

To examine the effect of AI adoption on productivity, we measure firm-level productivity using value-added total factor productivity (*Value-added TFP*). We focus on value-added TFP because it more closely captures underlying technical efficiency. Revenue-based TFP, by contrast, incorporates variation in output prices, demand conditions, and markups, thereby conflating production efficiency with firm-specific market power and demand shocks (Foster, Haltiwanger, and Syverson, 2008; De Loecker, 2011).

We obtain *Value-added TFP* as the residual from a log value-added Cobb–Douglas production function estimated with the generalized method of moments (GMM) approach of Wooldridge (2009).[16] The measurement of value added, capital, and labor inputs in the Cobb–Douglas production function follows Bennett, Stulz, and Wang (2020), and we use intermediate inputs, measured as the log of real materials, to proxy for unobserved productivity shocks as in Levinsohn and Petrin (2003). Real materials are constructed by deflating materials using the GDP implicit price deflator. Prior to estimation, we remove common industry-time variation by demeaning value added, intermediate inputs, capital, and labor using Fama–French 49 industry-year fixed effects, following İmrohoroğlu and Tüzel (2014). Firm-level variables, even after deflation, may still reflect differences in prices or markups rather than pure quantities. The inclusion of industry-year fixed effects absorbs components that are common within an industry-year, including shared movements in prices and demand. As a result, the estimation of firm-level TFP relies on within-industry-year variation. We then measure *Value-added TFP* using these residualized variables. Additional details are provided in Appendix Table B.1. As a widely used measure of productivity, value-added TFP can be interpreted as capturing differences in technological and organizational efficiency in transforming capital and labor into value added

[16] We adopt the generalized method of moments (GMM) framework of Wooldridge (2009), which implements the proxy-variable approaches of Levinsohn and Petrin (2003) as a system of moment conditions estimated jointly in a single step. This unified GMM formulation preserves the identification of the original Levinsohn–Petrin estimators and accommodates heteroskedasticity and serial correlation through a robust weighting matrix when estimating the production-function coefficients used to compute firm-level TFP. Our results are robust to alternative TFP estimates using the Levinsohn–Petrin method with the Ackerberg, Caves, and Frazer (2015) correction.

in the production process (Bennett, Stulz, and Wang, 2020).

We re-estimate the stacked DiD model with *Value-added TFP* as the dependent variable to compare changes in productivity for AI adopters before and after adoption relative to non-adopters. Columns (1)–(3) of Panel A in Table 8 show that the estimated coefficients on *AI* are positive and significant, indicating that AI adopters experience higher productivity relative to non-adopters, consistent with more efficient use of capital and labor in production. The impact of AI adoption on firm-level productivity is economically meaningful. The coefficient on *AI* in value-added TFP is 0.082 in year $t+1$, implying that adopters experience about 8.55% ($e^{0.082} - 1 \approx 0.0855$) higher value-added TFP one year after adoption relative to non-adopters.

To shed light on the value-added TFP response, we examine post-adoption effects on production inputs, focusing on capital and labor. These responses help assess whether higher productivity is accompanied by tangible-capital deepening or employment expansion, and thus whether it primarily reflects changes in the scale of production or mechanisms that rely more heavily on intangible and organizational complements. This distinction is particularly relevant to the innovation channel, through which AI can raise productivity by improving research efficiency and reshaping the innovation process. Because the key complements under this channel are largely intangible, productivity gains need not be accompanied by substantial increases in tangible capital, particularly in the near term (Cockburn, Henderson, and Stern, 2018; Brynjolfsson, Rock, and Syverson, 2021). The effect on labor demand is theoretically ambiguous. Labor demand may increase if AI and innovation complement human input in R&D and organizational implementation, but may decrease if AI substitutes for tasks previously performed by workers.

To test whether AI adoption affects labor input, we re-estimate Equation (5) with the log number of employees as the dependent variable. Columns (4)–(6) of Table 8 Panel A show that the coefficients on *AI* are insignificant, indicating no significant change in employment among AI adopters relative to non-adopters following adoption. Repeating the analysis for tangible capital yields similarly insignificant estimates, as shown in columns (7)–(9). Combined with

the evidence that AI adoption strengthens innovative output and efficiency, the increase in value-added TFP without significant growth in employment or tangible capital is consistent with the innovation channel, whereby AI raises the productivity of innovative inputs and these gains translate into higher firm-level productivity.[17]

These findings are consistent with endogenous growth models in which productivity growth arises from the production of new ideas through R&D activities (Romer, 1990; Aghion and Howitt, 1992). Within this framework, AI can raise productivity by improving the efficiency of the innovation process. The increases in inventor productivity and innovation efficiency documented in Table 5 Panel A provide evidence consistent with this mechanism. By enhancing the productivity of innovative effort, AI enables firms to generate and implement new ideas more efficiently, which can raise firm-level value-added TFP even without significant increases in physical capital or employment.

7.5 Contribution to Aggregate Productivity

Our results show that AI adopters experience higher value-added TFP following adoption than non-adopters. To gauge the potential aggregate effects, we map the estimated firm-level effects of AI adoption into an approximate contribution to aggregate value-added productivity using a first-order Hulten-style aggregation (Hulten, 1978). In the canonical Hulten-Domar framework, producer-level productivity shocks are aggregated using gross-output-based Domar weights. Because our firm-level productivity measure is value-added TFP, we implement the aggregation using firms' value-added shares, so the resulting measure should be interpreted as an approximation to aggregate value-added productivity that abstracts from intermediate-input propagation. A first-order approximation implies that the change in aggregate productivity can

[17] We also conduct cohort-specific event studies with lead–lag specifications to assess pre-adoption parallel trends and to trace the dynamics of post-adoption effects on productivity and other economic outcomes. As shown in Appendix Figure B.2, there is no evidence of pre-trends. Increases in *R&D*, *Tobin's q*, *Total q*, and *Value-added TFP* for adopters relative to non-adopters emerge only after AI adoption. Additionally, we implement a control function approach to address endogeneity concerns, using *Value-added TFP* and other economic outcomes as dependent variables. The results are robust, as reported in Appendix Table B.7 and Appendix Table B.8.

be written as the weighted sum of firm-level productivity changes:

$$d \ln A_h^{\text{agg}} = \sum_{i \in S} w_i \, d \ln A_{ih} + \sum_{j \in N} w_j \, d \ln A_{jh} \tag{10}$$

where S denotes AI adopters, N denotes non-adopters, $d \ln A_{ih}$ ($d \ln A_{jh}$) is the horizon-h change in firm-level value-added TFP for adopter i (non-adopter j), and w_i (w_j) denotes the firm's Domar-style weight, defined as its value-added share of GDP.

Let $w^S \equiv \sum_{i \in S} w_i$ and $w^N \equiv \sum_{j \in N} w_j$, and define the corresponding weighted-average value-added TFP changes for adopters and non-adopters as

$$g_h^S \equiv \frac{\sum_{i \in S} w_i d \ln A_{ih}}{w^S}, \qquad g_h^N \equiv \frac{\sum_{j \in N} w_j d \ln A_{jh}}{w^N}$$

Equation (10) can then be written as

$$d \ln A_h^{\text{agg}} = w^S g_h^S + w^N g_h^N \tag{11}$$

In our implementation, w^S is computed as AI adopters' aggregate value-added share of GDP in 2015 levels, thereby holding base-period shares fixed and abstracting from endogenous reallocation.[18] Empirically, the stacked DiD estimate $\hat{\beta}_h$ captures the average post-adoption change in value-added TFP for adopters relative to non-adopters at horizon h. We use $\hat{\beta}_h$ as an approximation to the difference in weighted-average value-added TFP changes between adopters and non-adopters:

$$\hat{\beta}_h \approx g_h^S - g_h^N \tag{12}$$

To capture potential spillover and competitive effects on non-adopters, we estimate an auxiliary

[18] In Domar (1961) and Hulten (1978), when intermediate inputs matter, the first-order effect of a producer-level productivity change on aggregate output is weighted by the producer's gross output relative to aggregate final demand, so the canonical Domar weights are sales based. As an implementation choice, we instead use value-added shares to align the aggregation with our firm-level measure of productivity. As a result, the aggregate contribution should be interpreted as an approximation to aggregate value-added productivity that abstracts from input-output propagation.

specification to obtain the estimate of $\hat{\beta}_h^N$, which we use to approximate the average change in value-added TFP among non-adopters:

$$\hat{\beta}_h^N \approx g_h^N \tag{13}$$

Substituting Equation (12) and Equation (13) into Equation (11) yields:

$$d \ln A_h^{\text{agg}} \approx w^S \hat{\beta}_h + (w^S + w^N)\hat{\beta}_h^N \tag{14}$$

When non-adopters do not experience significant changes in value-added TFP around the adoption year ($\hat{\beta}_h^N = 0$), Equation (14) simplifies to $d \ln A_h^{\text{agg}} \approx w^S \hat{\beta}_h$. More generally, when non-adopters are also affected, for example through spillover or competitive effects, the aggregate effect includes their contribution through the second term in Equation (14).

Table 8 Panel A shows that AI adopters experience an average increase in firm-level value-added TFP of 0.085 log points relative to non-adopters in a typical post-adoption year, where $0.085 = (0.082+0.08+0.093)/3$ is the average of estimated effects in years t+1, t+2, and t+3 after adoption. This corresponds to an increase of approximately 8.87% ($e^{0.085} - 1$). Given that AI adopters' value added accounts for 8.72% of GDP, their contribution to aggregate value-added TFP is approximately 0.0074 log points (0.0872×0.085), which is about 0.74% ($e^{0.0074} - 1$).

To assess whether non-adopters experience changes in value-added TFP around the adoption year, we re-estimate the stacked DiD specification by replacing the indicator AI with its components. Let *Post* equals 1 in the post-adoption year relative to a cohort's adoption year, and 0 otherwise. Let *Adopter* equals 1 for adopters in that cohort, and 0 otherwise. By construction, $AI \equiv Post \times Adopter$. We then estimate a stacked DiD regression of firm-level value-added TFP on *Post* and $Post \times Adopter$ with firm-by-cohort fixed effects that absorb all time-invariant firm heterogeneity within each cohort, including the effect of *Adopter*, and calendar year fixed effects that absorb shocks common to all firms in a given year. Standard errors are clustered at the firm level to account for serial correlation and repeated appearances of firms across cohorts. In this

specification, the coefficient on *Post* captures the average change in value-added TFP around adoption for non-adopters, while the coefficient on $Post \times Adopter$ captures the incremental change for AI-adopters relative to non-adopters.

As shown in Table 8 Panel B, the coefficients of *Post* are negative and statistically significant. The average of estimated effects on firm-level value-added TFP of non-adopter is -0.007 log points $(-(0.009+0.007+0.005)/3)$ in a typical post-adoption year. Since non-adopters account for 9.89% of GDP in 2015 levels, their direct contribution is approximately −0.0007 log points (−0.007×0.0989), corresponding to about −0.07%. This indicates a negative contribution of non-adopters to aggregate value-added TFP, consistent with spillover or competitive effects that partially offset aggregate gains from AI adoption. Applying Equation (14), the implied aggregate effect is 0.0061 log points (0.0872×0.085+(0.0872+0.0989)×(−0.007)). This corresponds to an increase of roughly 0.61% in aggregate value-added TFP in a typical post-adoption year, indicating that aggregate gains from AI adoption remain positive but are partially offset by negative effects on non-adopters. This estimate is on the conservative side, as it abstracts from input–output propagation and general equilibrium reallocation.

8 Conclusions

This paper investigates how AI adoption reshapes firm innovation and its implications for productivity. Using a novel firm-level measure that directly captures the timing of AI adoption and a stacked DiD design, we find that adopters generate more patents and their post-adoption patents receive more citations than those of non-adopters. AI adoption also changes innovation characteristics along three dimensions: novelty, direction, and reach. Channel tests suggest that AI transforms how firms innovate by improving the efficiency of innovative activity, changing the composition of inventor teams, and facilitating the recombination of specialized knowledge.

Following adoption, AI-adopting firms exhibit an approximately 8.87% increase in value-added TFP and higher investment in R&D, with little evidence of a corresponding expansion in tangible capital expenditure or a reduction in operating costs relative to non-adopters. This

pattern is more consistent with an innovation-oriented mechanism, in which AI raises research efficiency and reshapes the innovation process, than with immediate cost cutting or physical-capital driven expansion. Consistent with this interpretation, AI adopters experience higher firm valuations following adoption, suggesting that investors anticipate future benefits from AI adoption, including potential gains from innovation.

To gauge the potential aggregate implications of the firm-level TFP gains, we apply a Hulten-style aggregation and estimate a contribution of approximately 0.61% to aggregate value-added TFP growth in a representative post-adoption year. This implied aggregate contribution is a first-order, partial-equilibrium approximation that abstracts from the pace of diffusion and from general equilibrium forces such as input-output linkages and endogenous reallocation. To the extent that mechanisms such as intermediate-input propagation and factor reallocation amplify productivity gains, these estimates should be interpreted as conservative. Quantifying these forces is an important direction for future research.

REFERENCES

Acemoglu, D. 2025. The simple macroeconomics of AI. *Economic Policy* 40:13–58.

Acemoglu, D., D. Autor, J. Hazell, and P. Restrepo. 2022. Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics* 40:S293–340.

Acemoglu, D., and P. Restrepo. 2020. The wrong kind of AI? artificial intelligence and the future of labour demand. *Cambridge Journal of Regions, Economy and Society* 13:25–35.

Acharya, V. V., and Z. Xu. 2017. Financial dependence and innovation: The case of public versus private firms. *Journal of Financial Economics* 124:223–43.

Ackerberg, D. A., K. Caves, and G. Frazer. 2015. Identification properties of recent production function estimators. *Econometrica* 83:2411–51.

Aghion, P., U. Akcigit, A. Bergeaud, R. Blundell, and D. Hémous. 2019. Innovation and top income inequality. *The Review of Economic Studies* 86:1–45.

Aghion, P., and S. Bunel. 2024. AI and growth: Where do we stand? Policy Note.

Aghion, P., and P. Howitt. 1992. A model of growth through creative destruction. *Econometrica* 60:323–51.

Agrawal, A., J. Gans, and A. Goldfarb. 2018. Prediction, judgment, and complexity: A theory of decision-making and artificial intelligence. In *The Economics of Artificial Intelligence: An Agenda*, 89–110. University of Chicago Press.

Agrawal, A., J. McHale, and A. Oettl. 2019. Finding needles in haystacks: Artificial intelligence and recombinant growth. In *The Economics of Artificial Intelligence: An Agenda*, 149–74. University of Chicago Press.

———. 2024. Artificial intelligence and scientific discovery: A model of prioritized search. *Research Policy* 53:104989–.

Akcigit, U., S. Baslandze, and S. Stantcheva. 2016. Taxation and the international mobility of inventors. *American Economic Review* 106:2930–81.

Akcigit, U., M. A. Celik, and J. Greenwood. 2016. Buy, keep, or sell: Economic growth and the market for ideas. *Econometrica* 84:943–84.

Alderucci, D., S. Baviskar, L. Branstetter, N. Goldschlag, E. Hovy, A. Runge, P. Tambe, and N. Zolas. 2024. Quantifying the impact of AI on productivity and labor demand: Evidence from U.S. census microdata.

Alderucci, D., and D. Sicker. 2019. Applying artificial intelligence to the patent system. *Technology & Innovation* 20:415–25.

Almeida, H., and M. Campello. 2007. Financial constraints, asset tangibility, and corporate investment. *The Review of Financial Studies* 20:1429–60.

Ang, T. C., T. Chordia, V. V.-A. Mai, and H. Singh. 2022. The marketing capability premium. *Review of Asset Pricing Studies* 12:918–59.

Arts, S., B. Cassiman, and J. Hou. 2025. Technology differentiation, product market rivalry and ma transactions. *Strategic Management Journal* 46:837–62.

Azoulay, P., J. Krieger, and A. Nagaraj. 2025. Old moats for new models: Openness, control, and competition in generative artificial intelligence. *Entrepreneurship and Innovation Policy and the Economy* 4:7–46.

Babina, T. 2026. Understanding firms' AI efforts and their economic impact.

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

Baker, A. C., D. F. Larcker, and C. C. Wang. 2022. How much should we trust staggered difference-in-differences estimates? *Journal of Financial Economics* 144:370–95.

Banker, R. D., R. Huang, R. Natarajan, and S. Zhao. 2019. Market valuation of intangible asset: Evidence on SG&A expenditure. *The Accounting Review* 94:61–90.

Benner, M. J., and M. Tushman. 2002. Process management and technological innovation: A longitudinal study of the photography and paint industries. *Administrative Science Quarterly* 47:676–707.

Bennett, B., R. Stulz, and Z. Wang. 2020. Does the stock market make firms more productive? *Journal of Financial Economics* 136:281–306.

Bennett, N. R., J. L. Watson, R. J. Ragotte, A. J. Borst, D. L. See, C. Weidle, R. Biswas, Y. Yu, E. L. Shrock, R. Ault, et al. 2025. Atomically accurate de novo design of antibodies with RFdiffusion. *Nature* 1–11.

Bresnahan, T. F., E. Brynjolfsson, and L. M. Hitt. 2002. Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. *The Quarterly Journal of Economics* 117:339–76.

Bresnahan, T. F., and M. Trajtenberg. 2017. General purpose technologies 'engines of growth'? *The Quarterly Journal of Economics* 132:665–712.

Brynjolfsson, E., and L. M. Hitt. 2003. Computing productivity: Firm-level evidence. *Review of Economics and Statistics* 85:793–808.

Brynjolfsson, E., D. Li, and L. Raymond. 2025. Generative AI at work. *The Quarterly Journal of Economics* 140:889–942.

Brynjolfsson, E., and A. McAfee. 2017. What's driving the machine learning explosion? *Harvard Business Review* 18:118–.

Brynjolfsson, E., D. Rock, and C. Syverson. 2019. Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics. In *The Economics of Artificial Intelligence: An Agenda*, 23–60. University of Chicago Press.

———. 2021. The productivity J-Curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics* 13:333–72.

Cao, S., W. Jiang, B. Yang, and A. L. Zhang. 2023. How to talk when a machine is listening: Corporate disclosure in the age of AI. *The Review of Financial Studies* 36:3603–42.

Chen, M. A., and J. X. Wang. 2024. Displacement or augmentation? the effects of AI on workforce dynamics and firm value. *SSRN Working Paper* .

Chu, U. 2025. AI in R&D: Top use cases you need to know.

Cockburn, I. M., R. Henderson, and S. Stern. 2018. *The impact of artificial intelligence on innovation: An exploratory analysis*, vol. 24449. National Bureau of Economic Research Cambridge, MA, USA.

Cohn, J. B., Z. Liu, and M. I. Wardlaw. 2022. Count (and count-like) data in finance. *Journal of Financial Economics* 146:529–51.

Comment, R., and G. A. Jarrell. 1995. Corporate focus and stock returns. *Journal of Financial Economics* 37:67–87.

Corrado, C., C. Hulten, and D. Sichel. 2009. Intangible capital and US economic growth. *Review of Income and Wealth* 55:661–85.

Crouzet, N., J. C. Eberly, A. L. Eisfeldt, and D. Papanikolaou. 2022. The economics of intangible capital. *Journal of Economic Perspectives* 36:29–52.

De Loecker, J. 2011. Product differentiation, multiproduct firms, and estimating the impact of trade liberalization on productivity. *Econometrica* 79:1407–51.

Domar, E. D. 1961. On the measurement of technological change. *The Economic Journal* 71:709–29.

Eisfeldt, A. L., E. T. Kim, and D. Papanikolaou. 2022. Intangible value. *Critical Finance Review* 11:299–332.

Eisfeldt, A. L., G. Schubert, and M. B. Zhang. 2025. Generative AI and firm values. *Journal of Finance* Forthcoming.

Eloundou, T., S. Manning, P. Mishkin, and D. Rock. 2024. GPTs are GPTs: Labor market impact potential of llms. *Science* 384:1306–8.

Enache, L., and A. Srivastava. 2018. Should intangible investments be reported separately or commingled with operating expenses? new evidence. *Management Science* 64:3446–68.

Felten, E., M. Raj, and R. Seamans. 2021. Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal* 42:2195–217.

Fleming, L. 2001. Recombinant uncertainty in technological search. *Management Science* 47:117–32.

Foster, L., J. Haltiwanger, and C. Syverson. 2008. Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? *American Economic Review* 98:394–425.

Gao, H., P.-H. Hsu, and K. Li. 2018. Innovation strategy of private firms. *Journal of Financial and Quantitative Analysis* 53:1–32.

Gao, W., and J. Chou. 2015. Innovation efficiency, global diversification, and firm value. *Journal of Corporate Finance* 30:278–98.

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

Gormley, T. A., and D. A. Matsa. 2011. Growing out of trouble? corporate responses to liability risk. *The Review of Financial Studies* 24:2781–821.

———. 2016. Playing it safe? managerial preferences, risk, and agency conflicts. *Journal of Financial Economics* 122:431–55.

Griliches, Z. 1998. Patent statistics as economic indicators: A survey. In *R&D and Productivity: The Econometric Evidence*, 287–343. University of Chicago Press.

Hainmueller, J. 2012. Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies. *Political Analysis* 20:25–46.

Hall, B. H. 1990. The manufacturing sector master file: 1959-1987.

Hall, B. H., Z. Griliches, and J. A. Hausman. 1984. Patents and R&D: Is there a lag? *NBER Working Paper* .

Hall, B. H., A. Jaffe, and M. Trajtenberg. 2005. Market value and patent citations. *The RAND Journal of Economics* 36:16–38.

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

He, Q., and B. Qiu. 2025. Environmental enforcement actions and corporate green innovation. *Journal of Corporate Finance* 91:102711.

Hoberg, G., and V. Maksimovic. 2015. Redefining financial constraints: A text-based analysis. *The Review of Financial Studies* 28:1312–52.

Hoberg, G., and G. Phillips. 2016. Text-based network industries and endogenous product differentiation. *Journal of Political Economy* 124:1423–65.

Hulten, C. R. 1978. Growth accounting with intermediate inputs. *The Review of Economic Studies* 45:511–8.

İmrohoroğlu, A., and Ş. Tüzel. 2014. Firm-level productivity, risk, and return. *Management Science* 60:2073–90.

Jansen, J. J., F. A. Van Den Bosch, and H. W. Volberda. 2006. Exploratory innovation, exploitative innovation, and performance: Effects of organizational antecedents and environmental moderators. *Management Science* 52:1661–74.

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

Kogan, L., and D. Papanikolaou. 2019. Technological innovation, intangible capital, and asset prices. *Annual Review of Financial Economics* 11:221–42.

Kogan, L., D. Papanikolaou, A. Seru, and N. Stoffman. 2017. Technological innovation, resource allocation, and

growth. *The Quarterly Journal of Economics* 132:665–712.
Lee, P. 2010. Patent law and the two cultures. *Yale Law Journal* 120:2.
Lerner, J. 1994. The importance of patent scope: an empirical analysis. *The RAND Journal of Economics* 319–33.
Lerner, J., and A. Seru. 2022. The use and misuse of patent data: Issues for finance and beyond. *The Review of Financial Studies* 35:2667–704.
Levinsohn, J., and A. Petrin. 2003. Estimating production functions using inputs to control for unobservables. *The Review of Economic Studies* 70:317–41.
Li, D. 2011. Financial constraints, R&D investment, and stock returns. *The Review of Financial Studies* 24:2974–3007.
Li, K., and J. Wang. 2023. Inter-firm inventor collaboration and path-breaking innovation: Evidence from inventor teams post-merger. *Journal of Financial and Quantitative Analysis* 58:1144–71.
Lu, Y., G. M. Phillips, and J. Yang. 2024. The impact of cloud computing and AI on industry dynamics and concentration. *NBER Working Paper* .
Ludwig, J., and S. Mullainathan. 2024. Machine learning as a tool for hypothesis generation. *The Quarterly Journal of Economics* 139:751–827.
Ma, S., J. T. Tong, and W. Wang. 2022. Bankrupt innovative firms. *Management Science* 68:6971–92.
McElheran, K., M.-J. Yang, Z. Kroff, and E. Brynjolfsson. 2025. The rise of industrial AI in America: Microfoundations of the productivity J-Curve(s). Center for economic studies working paper, U.S. Census Bureau, Center for Economic Studies.
McPhee, D. 2024. AI algorithm identifies high-performing electrolytes for batteries.
Merchant, A., S. Batzner, S. S. Schoenholz, M. Aykol, G. Cheon, and E. D. Cubuk. 2023. Scaling deep learning for materials discovery. *Nature* 624:80–5.
Microsoft. 2024. Unlocking a New Era for Scientific Discovery with AI: How Microsoft's AI Screened Over 32 Million Candidates to Find a Better Battery.
Moretti, E. 2021. The effect of high-tech clusters on the productivity of top inventors. *American Economic Review* 111:3328–75.
Mühlroth, C., and M. Grottke. 2020. Artificial intelligence in innovation: How to spot emerging trends and technologies. *IEEE Transactions on Engineering Management* 69:493–510.
Ouellette, L. L., V. Fang, and N. T. Ouellette. 2025. How will AI affect patent disclosures? *Nature Biotechnology* 43:26–8.
Peters, R. H., and L. A. Taylor. 2017. Intangible capital and the investment-q relation. *Journal of Financial Economics* 123:251–72.
Raj, M., and R. Seamans. 2019. Artificial intelligence, labor, productivity, and the need for firm-level data. In A. Agrawal, J. Gans, and A. Goldfarb, eds., *The Economics of Artificial Intelligence: An Agenda*, chap. 22, 553–65. Chicago, IL: University of Chicago Press.
Rajan, R., H. Servaes, and L. Zingales. 2000. The cost of diversity: The diversification discount and inefficient investment. *The Journal of Finance* 55:35–80.
Romer, P. M. 1990. Endogenous technological change. *Journal of Political Economy* 98:S71–S102.
Setchi, R., I. Spasić, J. Morgan, C. Harrison, and R. Corken. 2021. Artificial intelligence for patent prior art searching. *World Patent Information* 64:102021–.
Stoffman, N., M. Woeppel, and M. D. Yavuz. 2022. Small innovators: No risk, no return. *Journal of Accounting and Economics* 74:101492–.
Trajtenberg, M., R. Henderson, and A. Jaffe. 1997. University versus corporate patents: A window on the basicness of invention. *Economics of Innovation and New Technology* 5:19–50.
Webb, M., N. Short, N. Bloom, and J. Lerner. 2018. Some facts of high-tech patenting.
Weitzman, M. L. 1998. Recombinant growth. *The Guarterly Journal of Economics* 113:331–60.
Whang, S. E., Y. Roh, H. Song, and J.-G. Lee. 2023. Data collection and quality challenges in deep learning: A data-centric AI perspective. *The VLDB Journal* 32:791–813.
Wooldridge, J. M. 2009. Estimating firm-level production functions using proxy variables to control for unobservables. *Economics Letters* 104:112–4.
———. 2015. Control function methods in applied econometrics. *Journal of Human Resources* 50:420–45.
Wuchty, S., B. F. Jones, and B. Uzzi. 2007. The increasing dominance of teams in production of knowledge. *Science* 316:1036–9.
Yoo, H. S., Y. L. Jung, and S.-P. Jun. 2023. Prediction of smes' R&D performances by machine learning for project selection. *Scientific Reports* 13:7598.

Yordy, L. R. 2021. The library of babel for prior art: Using artificial intelligence to mass produce prior art in patent law. *Vanderbilt Law Review* 74:521–62.

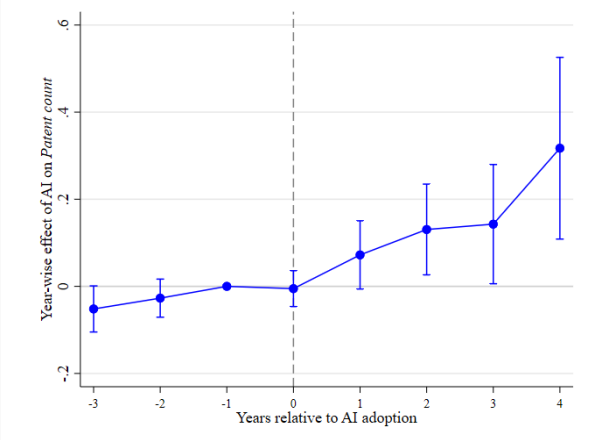
Zha, D., Z. P. Bhat, K.-H. Lai, F. Yang, Z. Jiang, S. Zhong, and X. Hu. 2025. Data-centric artificial intelligence: A survey. *ACM Computing Surveys* 57:1–42.

Zheng, X. 2025. How can innovation screening be improved? a machine learning analysis with economic consequences for firm performance. *Journal of Financial and Quantitative Analysis* 1–57.

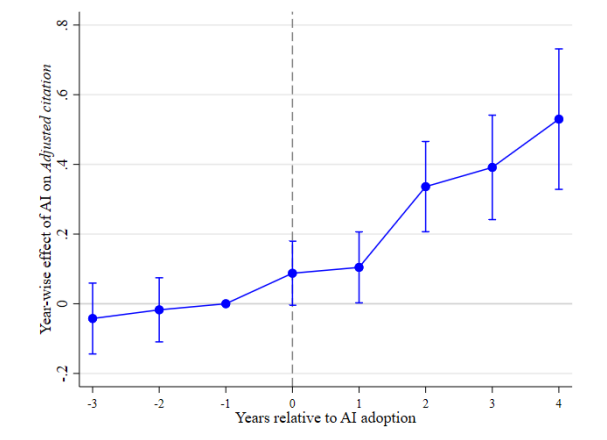
Figure 1
Dynamic Effects of AI Adoption

This figure presents coefficient estimates from the event-study specification examining the dynamic effects of AI adoption. The horizontal axis shows event time relative to AI adoption, with year −1 serving as the reference period. Error bars represent 95% confidence intervals. Panel A examines innovation quantity and quality: Figures (a) and (b) report estimates for *Patent count* and *Adjusted citation*, respectively. Panels B, C, and D examine innovation novelty, direction, and reach, respectively. Figures (c) and (d) report estimates for *Originality* and *Distance*; Figures (e) and (f) for *Exploitative* and *Explorative*; and Figures (g) and (h) for *Generality* and *Claim*. All specifications are estimated using Poisson pseudo-maximum likelihood (PPML) with year-cohort and firm-cohort fixed effects. Outcome variables are measured with a one-year lead. Standard errors are clustered at the firm level. Variable definitions are provided in Appendix Table B.1.

Panel A: Innovation Quantity and Quality



(a) *Patent count*

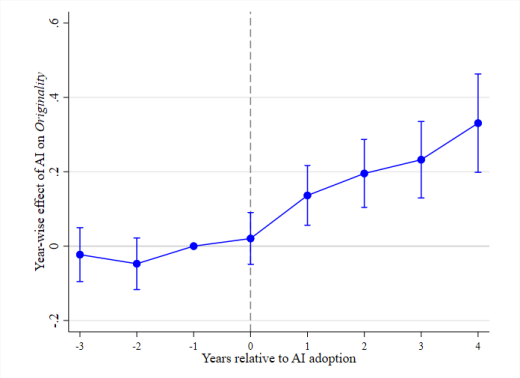


(b) *Adjusted citation*

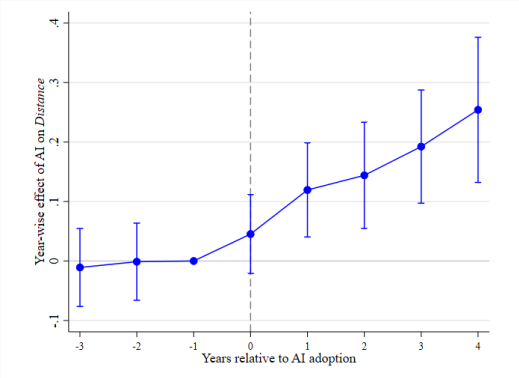
(continued on next page)

Figure 1 (continued)

Panel B: Innovation Novelty

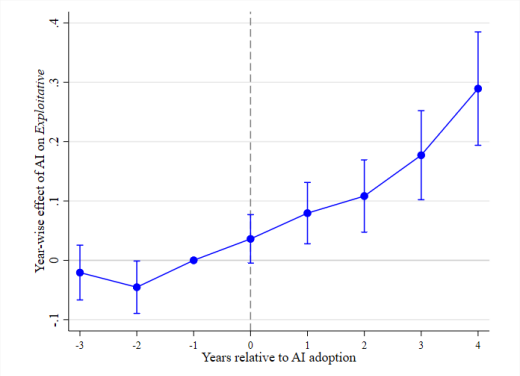


(c) *Originality*

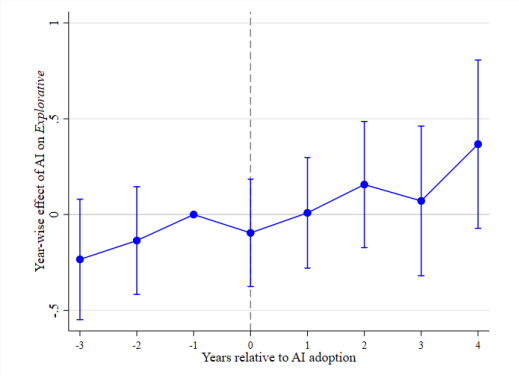


(d) *Distance*

Panel C: Innovation Direction

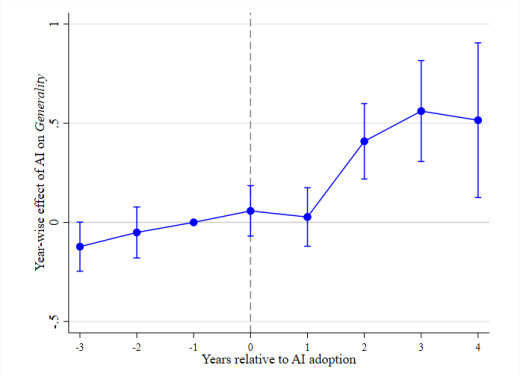


(e) *Exploitative*

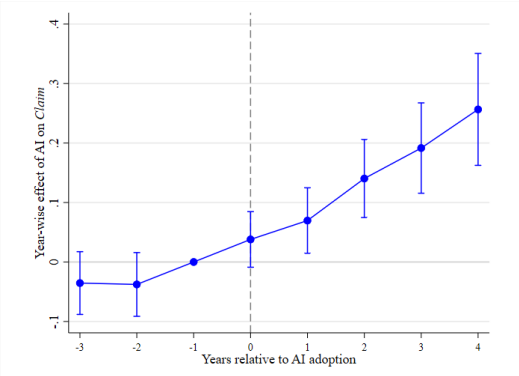


(f) *Explorative*

Panel D: Innovation Reach



(g) *Generality*



(h) *Claim*

Table 1
Summary Statistics

This table reports summary statistics for the firm-year panel from 2011 to 2021. *All firms* refers to the full sample of firms. *AI-adopting firms* are firms that adopt AI at any point during the sample period, while *Never-adopting firms* are firms that do not adopt AI during the sample period. Columns (1)-(3), (4)-(6), and (7)-(9) report summary statistics for all firms, AI-adopting firms, and never-adopting firms, respectively. Variable definitions are in Appendix Table B.1.

Variable	All firms			AI-adopting firms			Never-adopting firms		
	Mean	Median	S.D.	Mean	Median	S.D.	Mean	Median	S.D.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Patent count	24.685	2.000	74.205	65.971	9.000	123.191	9.308	1.000	32.161
Adjusted citation	0.521	0.248	0.753	0.672	0.529	0.760	0.465	0.127	0.743
Originality	0.271	0.265	0.254	0.291	0.306	0.223	0.264	0.233	0.264
Distance	0.242	0.143	0.269	0.295	0.266	0.258	0.222	0.068	0.271
Exploitative	0.499	0.636	0.458	0.643	0.929	0.438	0.446	0.400	0.454
Explorative	0.140	0.000	0.299	0.103	0.000	0.253	0.153	0.000	0.313
Generality	0.093	0.000	0.142	0.109	0.071	0.130	0.087	0.000	0.146
Claim	12.262	15.500	9.463	14.079	17.060	8.432	11.585	14.326	9.734
Value-added TFP	0.001	-0.049	0.670	-0.129	-0.173	0.606	0.051	0.004	0.686
Labor	0.791	0.924	2.075	2.358	2.405	1.649	0.192	0.289	1.901
Capital	5.741	5.750	2.461	5.111	5.095	2.284	6.669	6.819	2.025
Ln(total assets)	7.003	6.987	2.087	8.661	8.693	1.812	6.386	6.338	1.830
Cash	0.248	0.153	0.250	0.184	0.127	0.178	0.272	0.171	0.268
Tobin's q	2.414	1.762	1.855	2.481	1.856	1.797	2.390	1.721	1.876
Total q	1.716	0.911	3.075	1.849	1.097	2.671	1.667	0.838	3.212
R&D	0.078	0.021	0.127	0.047	0.017	0.071	0.090	0.023	0.141
Number of firms		2,600			518			2,082	

Table 2
AI and Innovation: Quantity and Quality

This table presents estimates from stacked DiD regressions of the effects of AI adoption on innovation outcomes. All specifications are estimated using Poisson pseudo-maximum likelihood (PPML). Panels A and B report estimates for patent quantity, measured by *Patent count*, and patent quality, measured by *Adjusted citation*, respectively. *AI* is an indicator equal to one if firm i in cohort c has adopted AI by year t, and zero otherwise. All specifications include year-cohort and firm-cohort fixed effects. All panels report regression estimates from $t+1$ to $t+3$, where t denotes the AI adoption year. Standard errors are clustered at the firm level. t-statistics are in parentheses. ***, ** and * indicate statistical significance at the 1%, 5%, and 10% level, respectively. Variable definitions are provided in Appendix Table B.1.

Panel A: Innovation Quantity			
	Patent count		
Variable	t+1	t+2	t+3
	(1)	(2)	(3)
AI	0.131***	0.149***	0.175***
	(3.02)	(3.60)	(4.56)
Year-Cohort FE	Yes	Yes	Yes
Firm-Cohort FE	Yes	Yes	Yes
Pseudo R^2	0.928	0.933	0.938
Observations	60,738	48,606	38,245
Panel B: Innovation Quality			
	Adjusted citation		
Variable	t+1	t+2	t+3
	(1)	(2)	(3)
AI	0.265***	0.277***	0.331***
	(6.05)	(5.71)	(5.96)
Year-Cohort FE	Yes	Yes	Yes
Firm-Cohort FE	Yes	Yes	Yes
Pseudo R^2	0.261	0.255	0.248
Observations	53,623	42,190	32,551

Table 3
AI and Innovation: Characteristics

This table presents estimates from stacked DiD regressions of the effects of AI adoption on innovation characteristics. Panels A, B, and C report estimates for innovation novelty, direction, and reach, respectively. Innovation novelty is measured by *Originality* and *Distance*, innovation direction by *Exploitative* and *Explorative*, and innovation reach by *Generality* and *Claim*. *AI* is an indicator variable equal to one if firm i in cohort c adopts AI in year t, and zero otherwise. All specifications include year-cohort and firm-cohort fixed effects and are estimated using PPML. All panels report regression estimates from $t+1$ to $t+3$, where t denotes the AI adoption year. Standard errors are clustered at the firm level. t-statistics are in parentheses. ***, ** and * indicate statistical significance at the 1%, 5%, and 10% level, respectively. Variable definitions are provided in Appendix Table B.1.

Panel A: Innovation Novelty						
	Originality			*Distance*		
Variable	t+1	t+2	t+3	t+1	t+2	t+3
	(1)	(2)	(3)	(4)	(5)	(6)
AI	0.169***	0.214***	0.195***	0.137***	0.162***	0.125***
	(5.55)	(6.52)	(5.20)	(4.54)	(4.98)	(3.51)
Year-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Pseudo R^2	0.0885	0.0846	0.0805	0.116	0.117	0.118
Observations	56,947	45,518	35,653	51,366	42,174	33,865
Panel B: Innovation Direction						
	Exploitative			*Explorative*		
Variable	t+1	t+2	t+3	t+1	t+2	t+3
	(1)	(2)	(3)	(4)	(5)	(6)
AI	0.147***	0.162***	0.154***	0.185**	0.204**	0.191*
	(6.50)	(6.59)	(5.69)	(2.12)	(1.97)	(1.69)
Year-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Pseudo R^2	0.109	0.104	0.0988	0.139	0.144	0.152
Observations	48,478	38,951	30,639	51,702	39,704	30,331
Panel C: Innovation Reach						
	Generality			*Claim*		
Variable	t+1	t+2	t+3	t+1	t+2	t+3
	(1)	(2)	(3)	(4)	(5)	(6)
AI	0.298***	0.317***	0.449***	0.142***	0.155***	0.161***
	(5.33)	(4.93)	(5.34)	(6.01)	(6.13)	(5.69)
Year-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Pseudo R^2	0.115	0.112	0.109	0.308	0.304	0.299
Observations	40,887	31,240	23,334	60,678	48,574	38,215

Table 4
The Impact of AI on Innovation: Control Function Approach

This table presents control function (CF) estimates of the effect of AI adoption on innovation using the stacked DiD sample. Column (1) reports the first-stage probit estimate for *AI exposure*. Columns (2)-(9) report second-stage estimates for innovation outcomes in year $t + 3$, where t denotes the AI adoption year. Columns (2) and (3) report estimates for innovation quantity and quality, respectively, while columns (4)-(9) report estimates for innovation characteristics. The second-stage models are estimated using PPML. *AI exposure* measures firm i's *ex ante* exposure to historically AI-strong universities through firm-university hiring networks as of 2010, and *Top 10 exposure* measures its exposure to top 10 universities ranked by the 2010 U.S. News & World Report through the same networks (Babina et al., 2024). *Residual* is the generalized residual obtained from the first-stage probit estimation. All specifications include year-cohort and industry-cohort fixed effects, with industries defined using the Fama-French 49-industry classification. Standard errors are computed using a firm-level cluster bootstrap with 1,000 replications. t-statistics are in parentheses. ***, ** and * indicate statistical significance at the 1%, 5%, and 10% level, respectively. Variable definitions are provided in Appendix Table B.1.

	1st Stage	2nd Stage							
Variable	*AI*	*Patent count*	*Adjusted citation*	*Originality*	*Distance*	*Exploitative*	*Explorative*	*Generality*	*Claim*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
AI exposure	0.572***								
	(4.68)								
AI		6.118***	1.485***	1.071***	0.832**	1.664***	-0.811	1.822**	1.348***
		(4.36)	(2.17)	(2.89)	(1.96)	(4.42)	(-0.93)	(2.03)	(3.68)
Top 10 exposure	-0.685	0.477	1.061**	0.274	0.367	0.388	0.730	1.266**	0.469**
	(-1.54)	(0.74)	(2.41)	(0.93)	(1.00)	(1.60)	(0.52)	(2.11)	(2.03)
Residual		-2.240***	-0.473	-0.394**	-0.194	-0.608***	0.311	-0.646	-0.479***
		(-2.95)	(-1.26)	(-1.99)	(-0.86)	(-3.00)	(0.65)	(-1.30)	(-2.61)
Year-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Wald χ^2	541.82								
Pseudo R^2	0.1149	0.427	0.0753	0.0464	0.0560	0.0565	0.0310	0.0466	0.0969
Observations	19,918	6,271	6,199	6,271	6,271	6,213	6,243	5,939	6,271

Table 5
The Impact of AI on Innovation: Channel

This table presents results from stacked DiD regressions examining the channels through which AI affects corporate innovation. Panel A reports results for the efficiency channel, measured by *Inventor productivity* and *Innovation efficiency*. Panel B presents results for the inventor composition channel, measured by *Team size* and *First-time inventor*. Panel C reports results for the knowledge recombination channel, measured by *Specialization* and *Core skill breadth*. *AI* is an indicator variable that equals 1 if firm i in cohort c adopts AI in year t, and 0 otherwise. All specifications include year-cohort and firm-cohort fixed effects and are estimated using PPML. All panels report regression estimates from $t + 1$ to $t + 3$, where t denotes the AI adoption year. Standard errors are clustered at the firm level. t-statistics are in parentheses. ***, ** and * indicate statistical significance at the 1%, 5%, and 10% level, respectively. Variable definitions and sources are provided in Appendix Table B.1.

Panel A: Efficiency Channel

	Inventor productivity			*Innovation efficiency*		
Variable	t+1	t+2	t+3	t+1	t+2	t+3
	(1)	(2)	(3)	(4)	(5)	(6)
AI	0.144***	0.166***	0.162***	0.161***	0.149***	0.166***
	(4.73)	(5.24)	(4.32)	(5.02)	(4.22)	(4.27)
Year-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Pseudo R^2	0.142	0.139	0.136	0.122	0.121	0.120
Observations	60,738	48,606	38,245	58,601	46,813	36,746

(continued on next page)

Table 5 (continued)

Panel B: Inventor Composition						
	Team size			First-time inventor		
Variable	t+1	t+2	t+3	t+1	t+2	t+3
	(1)	(2)	(3)	(4)	(5)	(6)
AI	0.142***	0.150***	0.154***	0.164***	0.214***	0.199***
	(5.69)	(5.43)	(5.08)	(2.93)	(3.52)	(2.80)
Year-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Pseudo R^2	0.237	0.231	0.227	0.0930	0.0932	0.0955
Observations	60,738	48,606	38,245	51,044	40,567	31,523
Panel C: Knowledge Recombination						
	Specialization			Core skill breadth		
Variable	t+1	t+2	t+3	t+1	t+2	t+3
	(1)	(2)	(3)	(4)	(5)	(6)
AI	0.148***	0.144***	0.148***	0.136***	0.144***	0.142***
	(6.37)	(5.32)	(5.02)	(5.75)	(5.44)	(5.00)
Year-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Pseudo R^2	0.0927	0.0886	0.0844	0.142	0.138	0.133
Observations	59,355	47,522	37,417	59,355	47,522	37,417

Table 6
The Impact of AI on Innovation: Heterogeneity

This table presents results from stacked DiDiD regressions examining cross-sectional heterogeneity in the impacts of AI on innovation. Panels A–D present differential effects by business focus, size, financial constraints, and technological capabilities, respectively. *AI* is an indicator variable that equals 1 if firm *i* in cohort *c* adopts AI in year *t*, and 0 otherwise. All specifications are estimated using PPML and include year-cohort and firm-cohort fixed effects, as well as interactions of these fixed effects with the corresponding subgroup indicator. All panels report regression estimates in year $t + 3$, where *t* denotes the AI adoption year. Standard errors are clustered at the firm level. *t*-statistics are in parentheses. ***, ** and * indicate statistical significance at the 1%, 5%, and 10% level, respectively. Variable definitions and sources are provided in Appendix Table B.1.

Panel A: Business Focus

Variable	*Patent count*	*Adjusted citation*	*Originality*	*Distance*	*Exploitative*	*Explorative*	*Generality*	*Claim*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AI	0.175***	0.315***	0.178***	0.119***	0.148***	0.144	0.420***	0.134***
	(4.41)	(5.52)	(4.50)	(3.20)	(5.08)	(1.21)	(4.70)	(4.52)
AI x *Focused firm*	0.445	0.368	0.411***	0.370**	0.251**	0.374	0.592	0.348***
	(1.63)	(1.16)	(2.98)	(2.12)	(2.37)	(0.88)	(1.37)	(2.84)
Year-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-Cohort-Focused firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort-Focused firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Pseudo R^2	0.938	0.251	0.0812	0.118	0.100	0.142	0.108	0.301
Observations	36,792	39,862	34,292	32,743	29,514	37,663	22,786	36,762

Panel B: Firm Size

Variable	*Patent count*	*Adjusted citation*	*Originality*	*Distance*	*Exploitative*	*Explorative*	*Generality*	*Claim*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AI	0.524**	-0.017	0.281**	-0.026	0.036	-0.057	0.071	0.057
	(2.08)	(-0.12)	(2.34)	(-0.21)	(0.37)	(-0.29)	(0.25)	(0.74)
AI x *Big firm*	-0.360	0.295*	-0.131	0.166	0.134	0.302	0.361	0.087
	(-1.41)	(1.83)	(-1.02)	(1.28)	(1.30)	(1.22)	(1.20)	(1.04)
Year-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-Cohort-Big firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort-Big firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Pseudo R^2	0.937	0.247	0.0811	0.117	0.101	0.150	0.109	0.302
Observations	36,215	30,979	33,729	32,171	29,009	28,977	22,333	36,185

(continued on next page)

Table 6 (continued)

Panel C: Financial Constraints								
Variable	*Patent count*	*Adjusted citation*	*Originality*	*Distance*	*Exploitative*	*Explorative*	*Generality*	*Claim*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AI	0.177***	0.374***	0.205***	0.151***	0.164***	0.208	0.510***	0.169***
	(3.73)	(5.93)	(4.94)	(3.70)	(5.26)	(1.54)	(5.00)	(5.18)
AI x *Unconstrained firm*	-0.018	-0.187	-0.026	-0.076	-0.024	-0.094	-0.256	-0.045
	(-0.24)	(-1.37)	(-0.28)	(-0.89)	(-0.36)	(-0.37)	(-1.38)	(-0.67)
Year-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-Cohort-Unconstrained firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort-Unconstrained firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Pseudo R^2	0.938	0.245	0.0811	0.118	0.100	0.150	0.108	0.301
Observations	36,792	31,509	34,292	32,726	29,514	29,465	22,786	36,762
Panel D: Technological Capabilities								
Variable	*Patent count*	*Adjusted citation*	*Originality*	*Distance*	*Exploitative*	*Explorative*	*Generality*	*Claim*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AI	0.138***	0.341***	0.219***	0.135***	0.190***	0.067	0.482***	0.170***
	(3.46)	(5.62)	(5.28)	(3.44)	(6.46)	(0.51)	(5.01)	(5.33)
AI x *High-tech firm*	0.149	-0.029	-0.094	-0.010	-0.165*	0.418	-0.296	-0.045
	(1.11)	(-0.19)	(-0.85)	(-0.09)	(-1.87)	(1.48)	(-1.28)	(-0.57)
Year-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-Cohort-High-tech firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort-High-tech firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Pseudo R^2	0.938	0.245	0.0812	0.118	0.100	0.150	0.108	0.301
Observations	36,792	31,509	34,292	32,743	29,514	29,465	22,786	36,762

Table 7
The Impact of AI on Economic Outcomes

This table presents results from stacked DiD regressions that estimate the effects of AI adoption on intangible and tangible investment, operating costs, and firm value. Panel A reports estimates for intangible and tangible investment, measured by *R&D* and *Capital expenditure*, respectively. Panel B reports estimates for operating costs, measured by *Production cost*, *Non-production cost*, and *Total operating cost*. Panel C reports estimates for firm value, measured by *Tobin's q* and *Total q*. All panels report linear regression estimates from $t+1$ to $t+3$, where t denotes the AI adoption year. *AI* is an indicator variable that equals one if firm i in cohort c adopts AI in year t, and zero otherwise. All specifications include year-cohort and firm-cohort fixed effects. Standard errors are clustered at the firm level. t-statistics are in parentheses. ***, ** and * indicate statistical significance at the 1%, 5%, and 10% level, respectively. Variable definitions and sources are provided in Appendix Table B.1.

Panel A: Intangible versus Tangible Investment

	R&D			*Capital expenditure*		
Variable	t+1	t+2	t+3	t+1	t+2	t+3
	(1)	(2)	(3)	(4)	(5)	(6)
AI	0.006***	0.006***	0.006***	-0.000	-0.001	-0.001
	(3.19)	(3.09)	(2.92)	(-0.12)	(-0.71)	(-0.58)
Year-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R^2	0.847	0.857	0.863	0.650	0.663	0.676
Observations	65,054	54,278	44,353	65,018	54,250	44,332

(continued on next page)

Table 7 (continued)

Panel B: Operating Costs

Variable	*Production cost*			*Non-production cost*			*Total operating cost*		
	t+1	t+2	t+3	t+1	t+2	t+3	t+1	t+2	t+3
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
AI	0.145	0.154	0.121	-0.001	-0.013	-0.017	0.132	0.063	0.000
	(1.04)	(1.14)	(1.01)	(-0.07)	(-0.91)	(-1.30)	(0.61)	(0.30)	(0.00)
Year-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R^2	0.617	0.619	0.639	0.761	0.787	0.816	0.623	0.632	0.661
Observations	65,054	54,278	44,353	56,546	47,539	39,179	65,054	54,278	44,353

Panel C: Firm Value

Variable	*Tobin's q*			*Total q*		
	t+1	t+2	t+3	t+1	t+2	t+3
	(1)	(2)	(3)	(4)	(5)	(6)
AI	0.273***	0.246***	0.163**	0.329***	0.337***	0.242***
	(4.27)	(3.59)	(2.40)	(3.72)	(3.59)	(2.63)
Year-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R^2	0.713	0.728	0.752	0.715	0.714	0.710
Observations	64,938	54,174	44,260	64,620	53,917	44,081

Table 8
The Impact of AI on Firm-Level Productivity

This table presents results from stacked DiD regressions that examine the impact of AI adoption on firm-level productivity, labor, and capital inputs. Panel A reports estimates comparing AI-adopting firms with non-adopting firms. Panel B examines whether non-adopters experience changes in these outcomes around cohort AI adoption years. All panels report linear regression estimates from $t+1$ to $t+3$, where t denotes the AI adoption year. *AI* is an indicator variable that equals one if firm i in cohort c adopts AI in year t, and zero otherwise. *Post* is an indicator variable that equals one in the post-adoption period relative to the adoption year of cohort c, and zero otherwise. *Adopter* is an indicator variable that equals one for adopters in cohort c, and zero otherwise. Panel A includes year-cohort and firm-cohort fixed effects, while Panel B includes year and firm-cohort fixed effects. Industries are defined using the Fama-French 49-industry classification. Standard errors are clustered at the firm level. t-statistics are in parentheses. ***, ** and * indicate statistical significance at the 1%, 5%, and 10% level, respectively. Variable definitions and sources are provided in Appendix Table B.1.

Panel A: Adopters vs. Non-Adopters

	Value-added TFP			*Labor*			*Capital*		
Variable	t+1	t+2	t+3	t+1	t+2	t+3	t+1	t+2	t+3
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
AI	0.082***	0.080***	0.093***	0.019	0.024	0.029	-0.009	0.001	0.007
	(4.73)	(4.44)	(4.67)	(0.86)	(1.11)	(1.39)	(-0.37)	(0.02)	(0.27)
Year-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R^2	0.878	0.887	0.893	0.977	0.981	0.983	0.979	0.982	0.984
Observations	59,981	49,987	40,841	59,981	49,987	40,841	59,981	49,987	40,841

Panel B: Non-adopter Post Effect

	Value-added TFP			*Labor*			*Capital*		
Variable	t+1	t+2	t+3	t+1	t+2	t+3	t+1	t+2	t+3
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Post	-0.009***	-0.007***	-0.005***	0.002	0.002	0.002	0.003	0.002	0.001
	(-5.59)	(-4.26)	(-3.48)	(0.96)	(1.17)	(1.14)	(1.52)	(0.87)	(0.64)
Adopter x *Post*	0.084***	0.083***	0.094***	0.020	0.025	0.030	-0.009	0.001	0.007
	(4.84)	(4.53)	(4.72)	(0.89)	(1.15)	(1.42)	(-0.36)	(0.05)	(0.28)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R^2	0.878	0.887	0.893	0.977	0.981	0.983	0.979	0.982	0.984
Observations	59,981	49,987	40,841	59,981	49,987	40,841	59,981	49,987	40,841

Appendix A. Robustness Checks

This section provides additional details on the robustness tests summarized in Section 4.6.

A.1 Entropy Balancing

One potential concern is that AI-adopting firms may differ systematically from non-adopting firms in observable characteristics, which could drive the observed effects. To mitigate this concern, we implement entropy balancing following Hainmueller (2012) to achieve covariate balance between treated (AI-adopting) and control (non-adopting) firms. Entropy balancing reweights the control group to match the treatment group on specified moments of selected covariates. Specifically, we balance on the means, variances, and skewnesses of *Ln(Total assets)*, *Cash*, and *Tobin's q*, using firm-level averages from financial years prior to 2015. The resulting weights are then applied to the full sample to ensure a balanced comparison across the treatment and control groups. Using the weighted sample, we re-estimate the baseline stacked DiD model (Equation (5)) and report the results in Panel A of Appendix Table B.4. The estimates are consistent with the baseline results, indicating that our findings are not driven by observable differences between AI-adopting and non-adopting firms.

A.2 Placebo Test

To further mitigate the concern that our results may be driven by pre-existing differences between AI-adopting and non-adopting firms, or some unobserved shocks coinciding with adoption that directly affect firms' innovation, we conduct a placebo test in which AI adoption is randomly assigned to firms that never adopt AI. For each year, we preserve the actual number of adopters observed in the original data and then randomly assign the same number of placebo adopters among the never-adopting firms. The assignment is performed without replacement, so each placebo-treated firm can be assigned at most one fictitious adoption year. We then construct a placebo treatment indicator based on these assignments and re-estimate Equation (5) using this placebo sample. Panel B of Appendix Table B.4 reports the estimation results. Across all specifications, the coefficients on AI are significantly negative or statistically insignificant. These results indicate that our results are unlikely to be driven by pre-existing differences between adopting and non-adopting firms or other confounding effects.

A.3 Alternative Control Group

In our baseline analysis, the control group includes both firms that never adopt AI technology (never-treated) and the pre-adoption observations of AI-adopting firms (not-yet-treated). A potential concern is that including not-yet-treated firms in the control group might introduce bias if their innovation behavior starts changing in anticipation of AI adoption. To assess the sensitivity of our results to the composition of the control group, we re-estimate the baseline stacked DiD model (Equation (5)) using only never-treated firms as the control group. Panel C of Appendix Table B.4 shows that the coefficients on AI are all positive and significant, which are similar to the

baseline estimates. This robustness check strengthens our identification strategy by showing that the baseline results are not driven by pre-treatment contamination or anticipatory behavior among future adopters.

A.4 Excluding AI Patents

The innovation measures in our baseline analysis are constructed from USPTO utility patents and include patents that explicitly utilise AI-related technologies. If AI adoption has a meaningful impact on firm-level innovation beyond AI-specific technologies, the estimated effects should remain when AI patents are excluded from the construction of these measures. To examine whether the effect of AI adoption is limited to AI-related invention, we exclude AI patents identified by Giczy, Pairolero, and Toole (2022) when constructing the innovation measures and re-estimate Equation (5). As shown in Panel D of Appendix Table B.4, the results are similar to our baseline estimates, supporting the interpretation that AI adoption enhances broader corporate innovation capabilities rather than merely stimulating activity within AI-specific technological domains.

Appendix B. Figures and Tables

Figure B.1
Diffusion of AI Adoption Across U.S. Geographies

This figure presents a heat map of the diffusion of AI adoption across the continental United States from 2016 to 2021. The color intensity reflects the number of AI-adopting firms within each U.S. Bureau of Economic Analysis (BEA) Economic Area in a given year, with darker colors indicating more adopters. Puerto Rico, Alaska, and Hawaii are excluded from the figure. Grey areas denote regions with no available data.

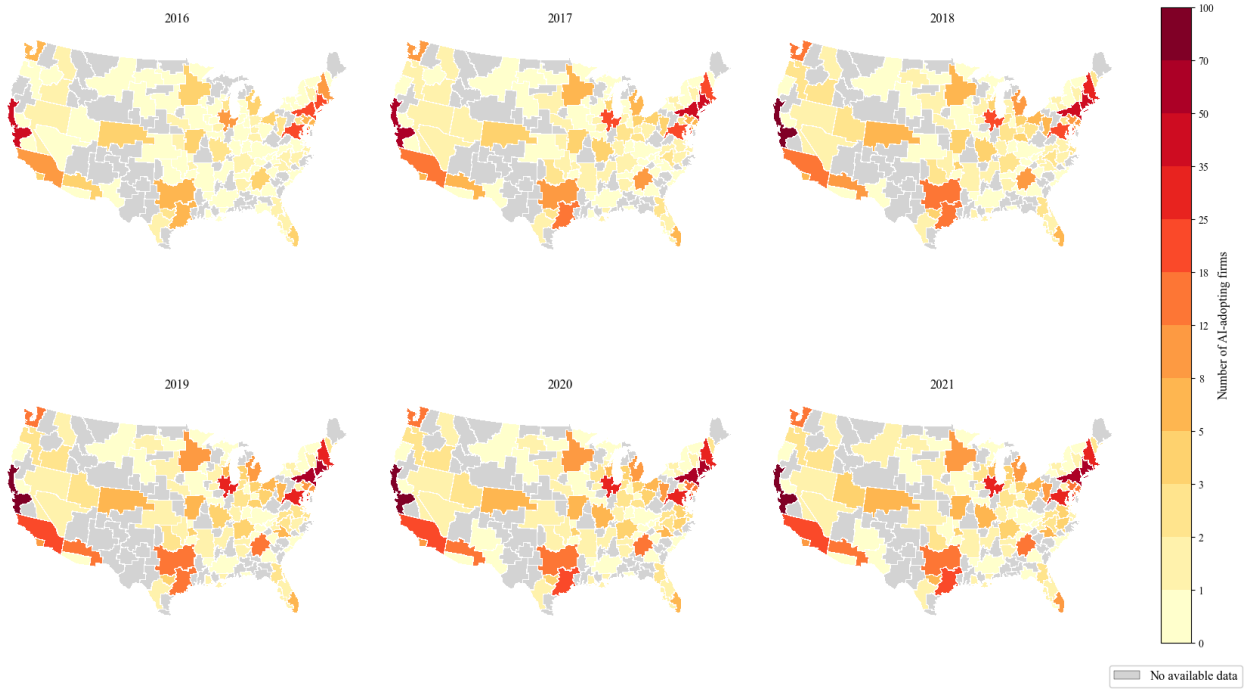
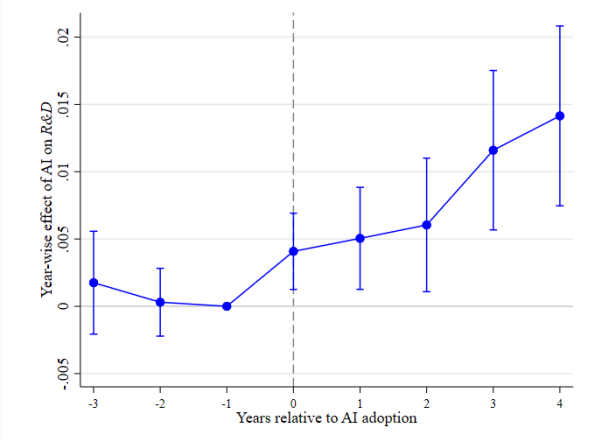


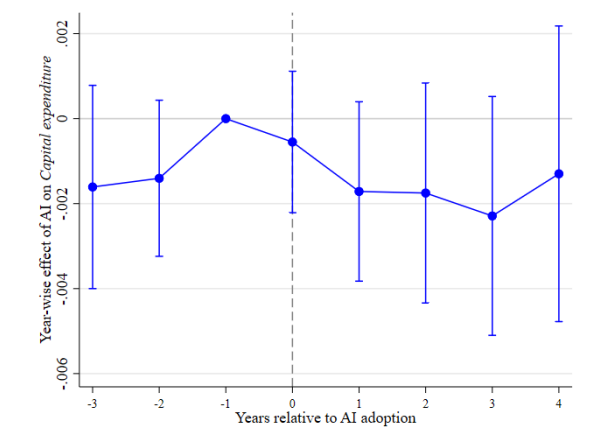
Figure B.2
Dynamic Effects of AI Adoption: Productivity and Economic Outcomes

This figure presents coefficient estimates from the event-study specification examining the dynamic effects of AI adoption on productivity and economic outcomes. The horizontal axis shows event time relative to AI adoption, with year −1 serving as the reference period. Error bars represent 95% confidence intervals. Panel A examines intangible versus tangible investment: Figures (a) and (b) report estimates for *R&D* and *Capital expenditure*, respectively. Panels B examines operating costs: Figures (c), (d), and (e) report estimates for *Production cost*, *Non-production cost*, and *Total operating cost*, respectively. Panel C examines firm value: Figures (f) and (g) report estimates for *Tobin's q* and *Total q*, respectively. Panel D examines firm-level productivity: Figures (h), (i), and (j) report estimates for *Value-added TFP*, *Labor*, and *Capital*, respectively. All specifications are estimated using PPML with year-cohort and firm-cohort fixed effects. Outcome variables are measured with a one-year lead. Standard errors are clustered at the firm level. Variable definitions are provided in Appendix Table B.1.

Panel A: Intangible versus Tangible Investment



(a) *R&D*

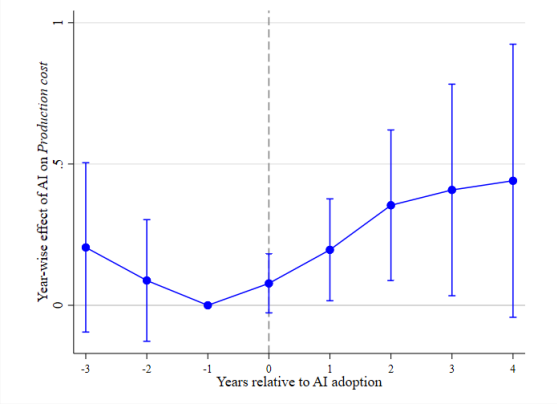


(b) *Capital expenditure*

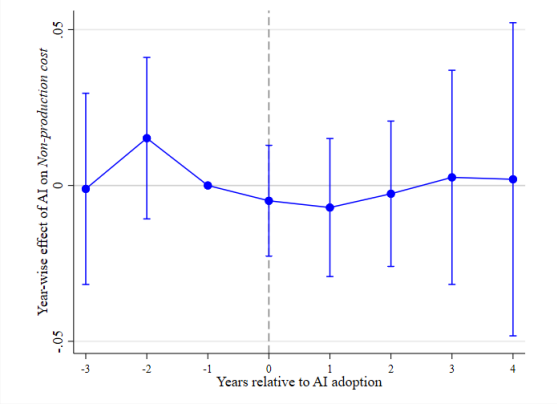
(continued on next page)

Figure B.2 (continued)

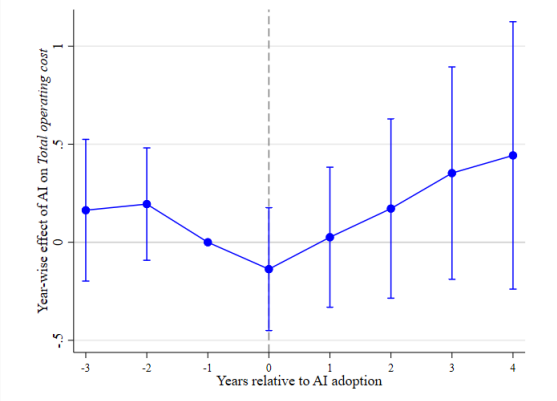
Panel B: Operating Costs



(c) *Production cost*

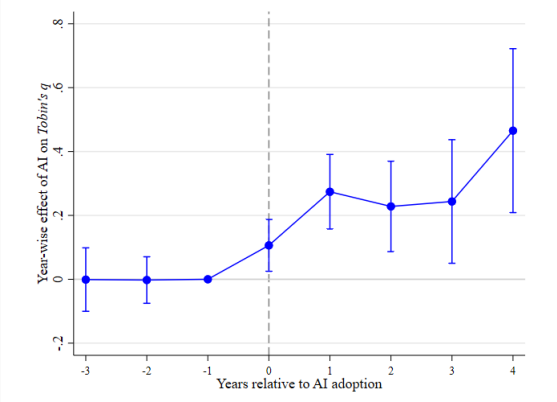


(d) *Non-production cost*

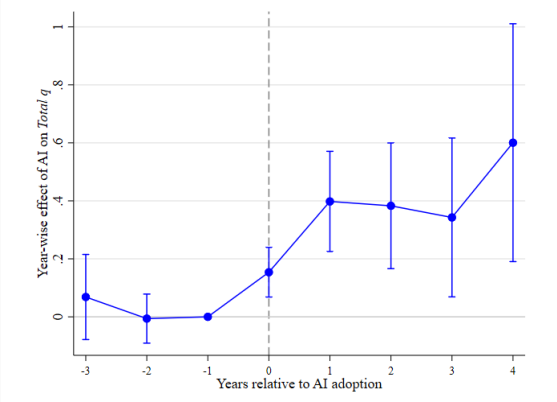


(e) *Total operating cost*

Panel C: Firm Value



(f) *Tobin's q*

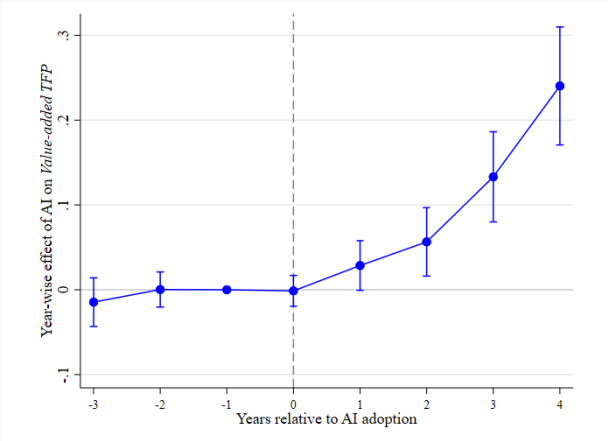


(g) *Total q*

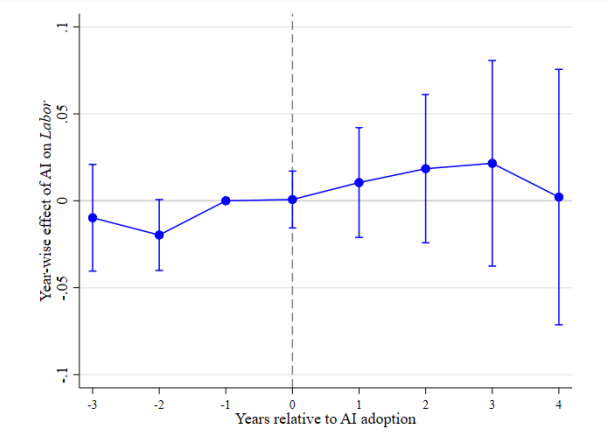
(continued on next page)

Figure B.2 (continued)

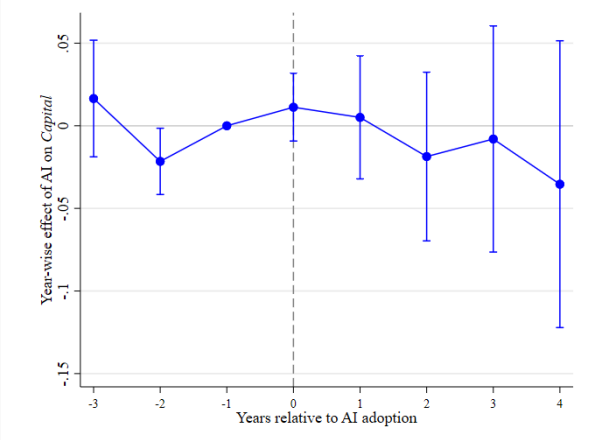
Panel D: Productivity



(h) *Value-added TFP*



(i) *Labor*



(j) *Capital*

Table B.1
Variable Definitions

Variable	Description	Source
	A. Innovation variables	
Patent count	The sum of granted patents filed by firm i in year t.	USPTO
Adjusted citation	The average of the time and technology class adjusted forward citations received by firm i to granted patents applied for in year t. Specifically, adjusted forward citations are the number of citations received by patent p applied by firm i in year t in class k scaled by the average number of citations received by all granted patents applied in year t in class k.	USPTO
Originality	One minus the average Herfindahl–Hirschman Index of the 2-digit Cooperative Patent Classification (CPC) technology classes in cited patents for granted patents filed by firm i in year t.	USPTO
Distance	The average technology distance between patent p and all patents previously filed by firm i up to year t.	USPTO
Exploitative	The proportion of exploitative patents filed by firm i in year t. Patent p is categorized as exploitative if 80% or more of its backward citations are based on firm i's existing expertise. Existing expertise is defined as the set of firm i's granted patents filed from the start of the fifth calendar year prior to the filing date of patent p up to the filing date of patent p, together with the patents cited by those granted patents during the same period.	USPTO
Explorative	The proportion of explorative patents filed by firm i in year t. Patent p is classified as explorative if 80% or more of its backward citations are not based on firm i's existing expertise. Existing expertise is defined as the set of firm i's granted patents filed from the start of the fifth calendar year prior to the filing date of patent p up to the filing date of patent p, together with the patents cited by those granted patents during the same period.	USPTO
Generality	One minus the average Herfindahl–Hirschman Index of the 2-digit Cooperative Patent Classification (CPC) technology classes in citing patents for granted patents filed by firm i in year t.	USPTO
Claim	The average number of claims made by granted patents filed by firm i in year t.	USPTO

(continued on next page)

Table B.1 (continued)

Variable	Description	Source
	B. AI variables	
AI	An indicator variable that equals 1 if firm i in cohort c has adopted AI technology by year t, and 0 otherwise.	SWZD
AI patent	Patent p filed by firm i in year t that is predicted by Giczy, Pairolero, and Toole (2022) to contain AI technology in any of the eight technology components (machine learning, natural language processing, computer vision, speech, knowledge processing, AI hardware, evolutionary computation, and planning and control) based on a 50% threshold.	Artificial Intelligence Patent Dataset
	C. Instrumental and related variables	
AI exposure	Firm i's exposure to AI-strong universities via the firm-university STEM worker hiring networks as of 2010.	Babina et al. (2024)
Top 10 exposure	Firm i's exposure to top 10 universities ranked by 2010 U.S. News & World Report via the firm-university STEM worker hiring networks as of 2010.	Babina et al. (2024)
	D. Financial variables	
Ln(Total assets)	The natural logarithm of the total assets (at).	CRSP/Compustat Merged
Cash	The sum of cash and short-term investments (che), divided by total assets (at).	CRSP/Compustat Merged
Tobin's q	Total assets (at) plus market value of equity (csho x prcc_f) minus book value of equity (ceq), divided by total assets (at). Missing prcc_f values are replaced with prcc_c.	CRSP/Compustat Merged
Total q	Market value of equity (csho x prcc_f) plus long-term debt (dltt) and debt in current liabilities (dlc), minus current assets (act), divided by the sum of total property, plant and equipment (ppegt) and replacement cost of intangible capital (Peters and Taylor, 2017). Replacement cost of intangible capital is the sum of a firm's externally purchased and internally created intangible capital. Externally purchased intangible capital is total intangible assets (intan), or 0 if the intangible assets are missing. A firm's internal intangible capital is the sum of knowledge capital and organization capital estimated using the perpetual inventory method. The stock of knowledge capital for firm i in year t is estimated by accumulating past R&D spending using the perpetual inventory method: $G_{i,t} = (1 - \delta_{R\&D})G_{i,t-1} + R\&D_{i,t}$, where $\delta_{R\&D}$ is the U.S. Bureau of Economic Analysis (BEA)'s industry-specific R&D depreciation rates and $R\&D_{i,t}$ is R&D expenses (xrd), or 0 if R&D expenses are missing. The stock of organization capital for firm i in year t is estimated by accumulating a fraction of past SG&A spending using the perpetual inventory method. SG&A is measured as selling, general and administrative expenses (xsga) minus the sum of R&D expenses (xrd) and in-process R&D expenses (rdip). Missing selling, general and administrative expenses (xsga), R&D expenses (xrd), and in-process R&D expenses (rdip) are replaced with 0. 30% of SG&A that is assumed to present an investment in intangible capital. The depreciation rate of SG&A is 20%.	Peters and Taylor (2017)

(continued on next page)

Table B.1 (continued)

Variable	Description	Source
	D. Financial variables (continued)	
Value-added TFP	Value-added total factor productivity. We estimate firm-level value-added TFP using a log Cobb–Douglas production function and the single-step GMM estimator of Wooldridge (2009). The baseline specification is: $y_{i,t} = \beta_0 + \beta_k k_{i,t} + \beta_l l_{i,t} + \epsilon_{i,t}$, where $y_{i,t}$ is the log of the value added of firm i in year t, $k_{i,t}$ is the log value of capital (K), $l_{i,t}$ is the log value of labor (L). We follow Bennett, Stulz, and Wang (2020) in defining value added, capital (K), and labor (L) variables. Value added is sales minus materials, deflated by the gross domestic product (GDP) implicit price deflator from Federal Reserve Economic Data (FRED). Sales are Compustat revenue (revt). Materials are computed as total expenses minus labor expenses. Total expenses are revenue (revt) less operating income before depreciation and amortization (oibdp). Labor expenses are non-missing wages (xlr). We use the non-missing wages (xlr) to calculate the average wage for each Fama-French 49 industry in year t. For firm-years with missing wages, we supplement missing xlr with industry-year average wage. Capital (K) is measured as gross plant, property, and equipment (ppegt) deflated by the price index for private fixed investment from Bureau of Economic Analysis (BEA), following the average age adjustment of İmrohoroğlu and Tüzel (2014), Hall (1990), and Brynjolfsson and Hitt (2003). Given that investment is made at different times in the past, we assume capital (K) was installed in a vintage year (calculated as year t minus the average age of capital; values larger or equal to 0.5 are rounded upward and values smaller than 0.5 are rounded downward), and apply the price index for private fixed investment in the corresponding vintage year as the deflator. We estimate the average age of capital for firm i in year t as accumulated depreciation (dpact; if dpact is missing, we use ppegt-ppent) divided by current depreciation (dp), then smooth this average age using a three-year moving average (or available years if fewer than three). Labor (L) is the number of employees (emp). We use intermediate inputs, measured as the log of real materials, as a proxy for unobserved productivity shocks following Levinsohn and Petrin (2003), implemented within the GMM framework of Wooldridge (2009). Real materials are constructed by deflating materials using the GDP implicit price deflator. Prior to estimation, we remove common industry-time variation by demeaning $y_{i,t}$, $k_{i,t}$, $l_{i,t}$, and the log of real materials using Fama–French 49 industry-year fixed effects, following İmrohoroğlu and Tüzel (2014). This transformation absorbs shocks that are common to firms within an industry-year, including shared movements in prices and demand. Firm-level TFP is thus estimated using within-industry-year variation. We then estimate the baseline specification using these residualized variables. *Value-added TFP* is computed as $y^{res}_{i,t} - \hat{\beta}_k k^{res}_{i,t} - \hat{\beta}_l l^{res}_{i,t}$, where $y^{res}_{i,t}$, $k^{res}_{i,t}$, $l^{res}_{i,t}$ denote the residuals from the industry-year fixed effects regressions.	CRSP/Compustat Merged, FRED, BEA

(continued on next page)

Table B.1 (continued)

Variable	Description	Source
	D. Financial variables (continued)	
Labor	The natural logarithm of labor (L). For details of labor (L) please see the definition of *Value-added TFP*.	CRSP/Compustat Merged
Capital	The natural logarithm of capital (K). For details of capital (K) please see the definition of *Value-added TFP*.	CRSP/Compustat Merged, BEA
R&D	R&D expenses (xrd) divided by total assets (at), or 0 if R&D expenses are negative or missing.	CRSP/Compustat Merged
Capital expenditure	Capital expenditures (capx) divided by total assets (at).	CRSP/Compustat Merged
Production cost	Cost of goods sold (cogs) divided by total sales (sale).	CRSP/Compustat Merged
Non-production cost	Selling, general and administrative expenses (xsga) minus R&D expenses (xrd) and advertising expenses (xad), divided by total sales (sale). Missing R&D expenses (xrd) and advertising expenses (xad) are replaced with 0.	CRSP/Compustat Merged
Total operating cost	Total operating expenses (xopr) divided by total sales (sale).	CRSP/Compustat Merged
	E. Channel variables	
Inventor productivity	The average number of granted patents filed by inventors of firm i in year t. The number of patents created by inventor m is calculated as the sum of all patents filed by inventor m in year t. For patents with multiple inventors, we follow Moretti (2021) to assign equally-weighted fractions of the patent to each of its inventors. We then aggregate the inventor–year patent counts to the firm–year level.	USPTO
Innovation efficiency	The ratio of the number of granted patents filed by firm i in year t to R&D expenditure in firm i in year t. The ratio is set to 0 if R&D expenses is reported as negative or missing (Gao and Chou, 2015).	USPTO
Team size	The average number of inventors per patent p filed by firm i in year t.	USPTO
First-time inventor	For firm i in year t, *first-time inventor* is the average (across the firm's patents filed in year t) of the patent-level share of inventors in each patent's inventor team who have no prior USPTO patent filings before year t.	USPTO

(continued on next page)

Table B.1 (continued)

Variable	Description	Source
	E. Channel variables (continued)	
Specialization	For firm *i* in year *t*, *Specialization* is the average (across the firm's patents filed in year *t*) of the patent-level average *Specialization* scores of inventors in each patent's inventor team. An inventor's *Specialization* score in year *t* is the Herfindahl–Hirschman Index computed from the distribution of her granted patents across the Cooperative Patent Classification (CPC) technology classes at the 3-digit level, using her granted patents applied for up to year *t−1* (Li and Wang, 2023). A higher *Specialization* score implies that the inventor is more specialized in her technology classes.	USPTO
Core skill breadth	For firm *i* in year *t*, *Core skill breadth* is the average (across the firm's patents filed in year *t*) of the patent-level number of distinct core skills possessed by inventors in each patent's inventor team. An inventor's core skill in year *t* is defined as the first three digits of the Cooperative Patent Classification (CPC) technology class in which she has the largest number of granted patents, based on her patent applications filed up to year *t-1* (Li and Wang, 2023).	USPTO
	F. Heterogeneity variables	
Focused firm	An indicator variable that equals 1 if firm *i* has only one business segment and no additional segments in 2015, and 0 otherwise.	Compustat Segments
Big firm	An indicator variable that equals 1 if the total assets of firm *i* in 2015 are equal to or larger than the industry average value in 2015, and 0 otherwise. Industries follow the Fama-French 49 Industry Classification.	CRSP/Compustat Merged
Unconstrained firm	An indicator variable that equals 1 if the risk of firm *i* to delay investments due to issues with liquidity is less than or equal to the 2015 industry average, and 0 otherwise. The risk of delaying investments due to liquidity issues is a text-based measure of financial constraints developed by Hoberg and Maksimovic (2015). Industries follow the Fama-French 49 Industry Classification.	Hoberg and Maksimovic (2015)
High-tech firm	An indicator variable that equals 1 if firm *i* operates in high-tech sectors in 2015, and 0 otherwise. We classify high-tech sectors as 2-digit NAICS sectors 51 ("Information") and 54 ("Professional, Scientific, and Technical Services") following Babina et al. (2024).	CRSP/Compustat Merged

Table B.2
Industry Distribution

This table presents the industry distribution of all firms in the firm-year panel sample from 2011 to 2021. *All firms* refers to the full sample of firms. *AI-adopting firms* are firms that adopt AI at any point during the sample period, while *Never-adopting firms* are firms that do not adopt AI during the sample period. Industries follow the Fama-French 49 Industry Classification.

	All firms	AI-adopting firms		Never-adopting firms	
Fama-French 49 Industry Classification	Count	Count	Percentage	Count	Percentage
Transportation	44	20	45.45	24	54.55
Entertainment	18	7	38.89	11	61.11
Retail	75	28	37.33	47	62.67
Utilities	51	18	35.29	33	64.71
Measuring and Control Equipment	63	22	34.92	41	65.08
Printing and Publishing	9	3	33.33	6	66.67
Aircraft	19	6	31.58	13	68.42
Wholesale	59	18	30.51	41	69.49
Business Services	95	28	29.47	67	70.53
Communication	62	18	29.03	44	70.97
Restaurants, Hotels, Motels	14	4	28.57	10	71.43
Tobacco Products	7	2	28.57	5	71.43
Computer Software	354	101	28.53	253	71.47
Healthcare	32	9	28.12	23	71.88
Consumer Goods	40	11	27.50	29	72.50
Personal Services	17	4	23.53	13	76.47
Computer Hardware	65	15	23.08	50	76.92
Petroleum and Natural Gas	79	18	22.78	61	77.22
Food Products	31	7	22.58	24	77.42
Recreation	23	5	21.74	18	78.26
Automobiles and Trucks	58	12	20.69	46	79.31
Apparel	20	4	20.00	16	80.00
Defense	10	2	20.00	8	80.00
Shipbuilding, Railroad Equipment	10	2	20.00	8	80.00
Candy and Soda	5	1	20.00	4	80.00
Electronic Equipment	218	42	19.27	176	80.73
Business Supplies	24	4	16.67	20	83.33
Beer and Liquor	6	1	16.67	5	83.33
Construction Materials	51	8	15.69	43	84.31
Steel Works Etc	26	4	15.38	22	84.62
Machinery	113	17	15.04	96	84.96
Rubber and Plastic Products	14	2	14.29	12	85.71
Chemicals	72	9	12.50	63	87.50
Medical Equipment	169	20	11.83	149	88.17
Almost Nothing	18	2	11.11	16	88.89
Shipping Containers	9	1	11.11	8	88.89
Electrical Equipment	55	6	10.91	49	89.09
Construction	27	2	7.41	25	92.59
Pharmaceutical Products	470	23	4.89	477	95.11
Agriculture	7	0	0.00	7	100.00
Coal	3	0	0.00	3	100.00
Fabricated Products	4	0	0.00	4	100.00
Non-Metallic and Industrial Metal Mining	10	0	0.00	10	100.00
Precious Metals	3	0	0.00	3	100.00
Textiles	4	0	0.00	4	100.00

Table B.3
Top 50 and AI-strong University List

This table presents the list of top 50 and AI-strong universities in the U.S. following Babina et al. (2024). *Top 50 university*, *Top 10 university*, and *Top 20 university* denote universities ranked among the top 50, top 10, and top 20 institutions by the U.S. News & World Report in 2010, respectively. *AI-strong university* refers to universities with strong AI research capabilities based on pre-2010 publications.

Top 50 university	Top 10 university	Top 20 university	AI-strong university
Amherst College	No	No	No
Boston College	No	No	No
Boston University	No	No	Yes
Bowdoin College	No	No	No
Brandeis University	No	No	No
Brown University	No	Yes	Yes
California Institute of Technology	Yes	Yes	Yes
Carleton College	No	No	No
Carnegie Mellon University	No	No	Yes
Case Western Reserve University	No	No	No
Claremont McKenna College	No	No	No
College of William and Mary all campuses	No	No	No
Columbia University in the City of New York	Yes	Yes	No
Cornell University all campuses	No	Yes	Yes
Dartmouth College	No	Yes	No
Duke University	Yes	Yes	No
Emory University	No	Yes	No
Georgetown University	No	No	No
Georgia Institute of Technology all campuses	No	No	Yes
Harvard University	Yes	Yes	No
Johns Hopkins University	No	Yes	Yes
Lehigh University	No	No	No
Massachusetts Institute of Technology	Yes	Yes	Yes
Middlebury College	No	No	No
New York University	No	No	Yes
Northeastern University	No	No	Yes
Northwestern University	No	Yes	Yes
Ohio State University all campuses	No	No	Yes
Pepperdine University	No	No	No
Pomona College	No	No	No
Princeton University	Yes	Yes	Yes
Rensselaer Polytechnic Institute	No	No	Yes
Rice University	No	Yes	No
Stanford University	Yes	Yes	Yes
Swarthmore College	No	No	No
Tufts University	No	No	No
Tulane University	No	No	No
University of California, Berkeley	No	Yes	Yes
University of California, Davis	No	No	No
University of California, Irvine	No	No	No
University of California, Los Angeles	No	Yes	Yes
University of California, San Diego	No	No	Yes
University of California, Santa Barbara	No	No	Yes
University of Chicago	Yes	Yes	No
University of Florida	No	No	Yes
University of Georgia	No	No	No
University of Illinois at Urbana-Champaign	No	No	Yes
University of Michigan all campuses	No	No	Yes
University of North Carolina Chapel Hill	No	No	No
University of Notre Dame	No	Yes	No
University of Pennsylvania	Yes	Yes	Yes
University of Rochester	No	No	No
University of Southern California	No	No	Yes
University of Texas Austin	No	No	Yes
University of Virginia all campuses	No	No	No
University of Wisconsin Madison	No	No	No
Vanderbilt University	No	Yes	No
Villanova University	No	No	No
Wake Forest University	No	No	No
Washington University in St. Louis	No	Yes	No
Washington and Lee University	No	No	No
Wellesley College	No	No	No
Williams College	No	No	No
Yale University	Yes	Yes	No

Table B.4
AI and Innovation: Robustness Tests

This table presents the results of stacked DiD regressions on the effect of AI on innovation in different robustness tests. Panel A reports estimates using an entropy-balanced sample. The firm-year panel sample is re-weighted based on the averages of *Ln(total assets)*, *Cash*, and *Tobin's q* in financial years before 2015 so that the means, variances, and skewnesses of the control group match the values of the same moments in the treated group between 2011 to 2015. Panel B reports the results of a placebo test in which AI adoption is randomly assigned to firms that never adopt AI. For each year, we preserve the actual number of adopters observed in the original data, switch off the treatment status from true adopters, and then randomly assign the same number of "placebo adopters'' among the never-adopting firms. The assignment is performed without replacement, ensuring that each placebo-treated firm can be assigned an adoption year at most once. Panel C presents estimates using the never-treated firms (i.e., firms that never adopt AI technology during the sample period) as the control group. Panel D reports results of the effect of AI on non-AI innovation. We exclude AI patents identified by Giczy, Pairolero, and Toole (2022) from the patent-level sample and re-construct corporate innovation measures without AI patents. In all panels, columns (1) and (2) report estimates for innovation quantity and quality, respectively, while columns (3)-(8) report estimates for innovation characteristics. *AI* is an indicator variable that equals 1 if firm i in cohort c uses AI technology in year t, and 0 otherwise. All specifications include year-cohort and firm-cohort fixed effects and are estimated using PPML. All panels report regression estimates in year $t + 3$, where t denotes the AI adoption year. Standard errors are clustered at the firm level. t-statistics are in parentheses. ***, ** and * indicate statistical significance at the 1%, 5%, and 10% level, respectively. Variable definitions and sources are in Appendix Table B.1.

Panel A: Entropy Balancing

Variable	*Patent count*	*Adjusted citation*	*Originality*	*Distance*	*Exploitative*	*Explorative*	*Generality*	*Claim*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AI	0.129*	0.153**	0.111***	0.043	0.083***	0.183	0.240**	0.100***
	(1.87)	(2.18)	(2.70)	(0.96)	(2.75)	(1.35)	(2.23)	(3.17)
Year-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Pseudo R-squared	0.944	0.209	0.0710	0.100	0.0781	0.181	0.0874	0.291
Observations	36,096	30,860	33,685	32,273	28,987	29,025	22,386	36,066

(continued on next page)

Table B.4 (continued)

Panel B: Placebo Test								
Variable	*Patent count*	*Adjusted citation*	*Originality*	*Distance*	*Exploitative*	*Explorative*	*Generality*	*Claim*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AI	-0.211***	-0.274***	-0.086	-0.069	-0.123***	0.084	-0.348**	-0.084*
	(-2.83)	(-3.12)	(-1.57)	(-1.26)	(-2.61)	(0.63)	(-2.50)	(-1.85)
Year-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Pseudo R^2	0.945	0.229	0.077	0.114	0.091	0.165	0.100	0.305
Observations	34,993	30,227	32,672	31,273	28,429	28,937	22,629	34,976
Panel C: Never-Treated Firms as the Control Group								
Variable	*Patent count*	*Adjusted citation*	*Originality*	*Distance*	*Exploitative*	*Explorative*	*Generality*	*Claim*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AI	0.253***	0.376***	0.225***	0.163***	0.189***	0.201*	0.500***	0.187***
	(4.56)	(6.17)	(5.66)	(4.29)	(6.28)	(1.74)	(5.42)	(6.16)
Year-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Pseudo R^2	0.918	0.257	0.0852	0.121	0.104	0.138	0.119	0.300
Observations	31,495	26,424	29,313	27,505	24,854	24,607	18,202	31,465
Panel D: Excluding AI Patents								
Variable	*Patent count*	*Adjusted citation*	*Originality*	*Distance*	*Exploitative*	*Explorative*	*Generality*	*Claim*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AI	0.212***	0.340***	0.186***	0.119***	0.150***	0.295**	0.448***	0.169***
	(4.86)	(6.25)	(4.95)	(3.37)	(5.62)	(2.44)	(5.30)	(6.05)
Year-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Pseudo R^2	0.930	0.245	0.0804	0.116	0.0995	0.151	0.116	0.298
Observations	37,694	31,747	34,882	33,117	30,268	29,036	22,354	37,672

Table B.5
AI and Innovation: Non-High-Tech Firms

This table presents the results of stacked DiD regressions on the effect of AI on innovation for non-high-tech firms. We exclude firms that operate in high-tech sectors in 2015. We classify high-tech sectors as 2-digit NAICS sectors 51 ("Information") and 54 ("Professional, Scientific, and Technical Services") following Babina et al. (2024). We present PPML results in year $t + 3$ following AI adoption in year t. Columns (1) and (2) report estimates for innovation quantity and quality, respectively, while columns (3)-(8) report estimates for innovation characteristics. *AI* is an indicator variable that equals 1 if firm i in cohort c uses AI technology in year t, and 0 otherwise. All specifications include Year-Cohort FE and Firm-Cohort FE. Standard errors are clustered at the firm level. t-statistics are in parentheses. ***, ** and * indicate statistical significance at the 1%, 5%, and 10% level, respectively. Variable definitions and sources are provided in Appendix Table B.1.

Variable	*Patent count*	*Adjusted citation*	*Originality*	*Distance*	*Exploitative*	*Explorative*	*Generality*	*Claim*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AI	0.138***	0.341***	0.219***	0.135***	0.190***	0.067	0.482***	0.170***
	(3.46)	(5.62)	(5.28)	(3.44)	(6.46)	(0.51)	(5.01)	(5.32)
Year-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Pseudo R^2	0.938	0.249	0.0808	0.115	0.103	0.147	0.111	0.299
Observations	31,679	27,109	29,771	28,344	25,535	25,215	19,381	31,666

Table B.6
The Heterogeneous Impact of AI on Innovation: Product Market Competition

This table presents results from stacked DiDiD regressions examining cross-sectional heterogeneity in the impacts of AI on corporate innovation across firms with different product market competition. We present PPML results in year $t + 3$ following AI adoption in year t. Columns (1) and (2) report estimates for innovation quantity and quality, respectively, while columns (3)-(8) report estimates for innovation characteristics. *AI* is an indicator variable that equals 1 if firm i in cohort c uses AI technology in year t, and 0 otherwise. *High competition* is an indicator variable that equals 1 if firm i's TNIC-based Herfindahl-Hirschman Index (HHI) is equal to or below the 2015 average, and 0 otherwise. The TNIC-based HHI is a text-based measure of product market competition developed by Hoberg and Phillips (2016). All specifications include Year–Cohort and Firm–Cohort fixed effects, as well as subgroup-specific Year–Cohort and Firm–Cohort fixed effects obtained by interacting these fixed effects with the indicator variable for product market competition. Standard errors are clustered at the firm level. t-statistics are in parentheses. ***, ** and * indicate statistical significance at the 1%, 5%, and 10% level, respectively. Variable definitions and sources are provided in Appendix Table B.1.

Variable	*Patent count*	*Adjusted citation*	*Originality*	*Distance*	*Exploitative*	*Explorative*	*Generality*	*Claim*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AI	0.180***	0.358***	0.104*	0.083	0.077	0.212	0.518***	0.077*
	(3.21)	(4.19)	(1.78)	(1.54)	(1.62)	(1.12)	(3.44)	(1.87)
AI x *High competition*	-0.011	-0.071	0.155**	0.083	0.120**	-0.018	-0.128	0.095
	(-0.15)	(-0.63)	(2.02)	(1.14)	(2.04)	(-0.07)	(-0.70)	(1.57)
Year-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-Cohort-High competition FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort-High competition FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Pseudo R^2	0.938	0.251	0.0811	0.118	0.100	0.150	0.108	0.307
Observations	36,792	39,862	34,292	32,743	29,514	29,465	22,786	36,762

Table B.7
The Impact of AI on Firm-Level Productivity: Control Function Approach

This table presents control function (CF) estimates of the effect of AI adoption on on productivity, labor and capital using the stacked DiD sample. Column (1) reports the first-stage probit estimate for *AI exposure*. Columns (2)-(4) report second-stage estimates for productivity, labor, and capital in year $t + 3$, where t denotes the AI adoption year. Columns (2) reports estimates for *Value-added TFP*, while columns (3)-(4) report estimates for *Labor* and *Capital*, respectively. The second-stage models are estimated using linear regressions. *AI exposure* measures firm i's *ex ante* exposure to historically AI-strong universities through firm-university hiring networks as of 2010, and *Top 10 exposure* measures its exposure to top 10 universities ranked by the 2010 U.S. News & World Report through the same networks (Babina et al., 2024). *Residual* is the generalized residual obtained from the first-stage probit estimation. All specifications include year-cohort and industry-cohort fixed effects, with industries defined using the Fama-French 49-industry classification. Standard errors are computed using a firm-level cluster bootstrap with 1,000 replications. t-statistics are in parentheses. ***, ** and * indicate statistical significance at the 1%, 5%, and 10% level, respectively. Variable definitions are provided in Appendix Table B.1.

	1st Stage	2nd Stage		
Variable	*AI*	*Value-added TFP*	*Labor*	*Capital*
	(1)	(2)	(3)	(4)
AI exposure	0.572***			
	(4.68)			
AI		0.692**	0.836	1.337
		(2.00)	(0.97)	(1.31)
Top 10 exposure	-0.685	0.114	-1.299*	-1.141
	(-1.54)	(0.27)	(-1.86)	(-1.51)
Residual		-0.418**	0.589	0.397
		(-2.21)	(1.24)	(0.71)
Year-Cohort FE	Yes	Yes	Yes	Yes
Industry-Cohort FE	Yes	Yes	Yes	Yes
Wald χ^2	541.82			
Pseudo R^2	0.1149			
Adjusted R^2		0.0121	0.396	0.443
Observations	19,918	6,089	6,089	6,089

Table B.8
The Impact of AI on Economic Outcomes: Control Function Approach

This table presents the results of Control Function (CF) regressions on intangible versus tangible investment, operating costs, and firm value using the stacked DiD sample. Column (1) reports the first-stage probit estimate for *AI exposure*. Columns (2)-(8) report second-stage estimates for economic outcomes in year $t + 3$, where t denotes the AI adoption year. Columns (2)-(3) presents results on intangible versus tangible investment, columns (4)-(6) presents results on operating costs, and columns (7)-(8) present results on firm value. The second-stage models are estimated using linear regressions. *AI exposure* is firm i's *ex ante* exposure to historically AI-strong universities via the firm-university STEM worker hiring networks as of 2010 (Babina et al., 2024). *Top 10 exposure* is firm i's exposure to top 10 universities ranked by 2010 U.S. News & World Report via the firm-university STEM worker hiring networks as of 2010 (Babina et al., 2024). *Residual* is the regression residuals from the first-stage estimation. All regressions include Year-Cohort FE and Industry-Cohort FE. Industries follow the Fama-French 49 Industry Classification. Standard errors are computed using a firm-level cluster bootstrap with 1,000 replications. t-statistics are in parentheses. ***, ** and * indicate statistical significance at the 1%, 5%, and 10% level, respectively. Variable definitions and sources are provided in Appendix Table B.1.

	1st Stage	2nd Stage						
Variable	*AI*	*R&D*	*Capital expenditure*	*Production cost*	*Non-production cost*	*Total operating cost*	*Tobin's q*	*Total q*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AI exposure	0.572***							
	(4.68)							
AI		0.137***	-0.012	1.281	0.149	1.002	3.997***	5.276**
		(4.10)	(-0.97)	(1.06)	(0.54)	(0.54)	(3.76)	(2.48)
Top 10 exposure	-0.685	0.173**	-0.014**	9.213	0.802	11.291	1.760	2.326
	(-1.54)	(2.54)	(-2.34)	(1.54)	(0.88)	(1.47)	(1.61)	(1.77)
Residual		-0.085***	0.006	-0.916	-0.167	-0.952	-2.137***	-2.881**
		(-4.55)	(0.97)	(-1.38)	(-1.12)	(-0.93)	(-3.62)	(-2.36)
Year-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Wald χ^2	541.82							
Pseudo R^2	0.1149							
Adjusted R^2		0.432	0.190	0.102	0.104	0.0938	0.158	0.0911
Observations	19,918	6,284	6,283	6,284	5,503	6,284	6,265	6,239