Artificial Intelligence and Innovation: Quantity, Quality, and Strategy*

Viral V. Acharya Xuehua Shao Zhaoxia Xu

Abstract

This paper examines how firms' adoption of artificial intelligence (AI) affects innovation. Using a novel firm-level usage-based measure of AI adoption and a stacked difference-in-differences design that exploits staggered adoption, we find that adopters subsequently increase patenting relative to non-adopters. Their post-adoption patents receive more citations, contain more claims, and include a higher share of exploitative patents that build on their existing technologies, and exhibit greater originality and generality and span more technologically distant classes. These patterns indicate that AI facilitates both refinement of existing knowledge and exploration of new areas. The results are robust to a control function approach that addresses endogenous adoption. These effects are stronger for firms with a more focused business scope, consistent with high-quality coherent data enhancing the productivity of AI, and they are not concentrated among larger or financially unconstrained firms. Evidence on mechanisms is consistent with AI improving efficiency, reshaping inventor composition, and enhancing knowledge recombination.

Keywords: Artificial Intelligence, Innovation, Patents

*Acharya is at New York University Stern School of Business (vva1@stern.nyu.edu). Shao (xuehua.shao@unsw.edu.au) and Xu (zhaoxia.xu@unsw.edu.au) are at University of New South Wales.

1 Introduction

Recent advances in artificial intelligence (AI) are reshaping firm behavior. Firms increasingly deploy AI in the innovation process, including idea generation and screening, forecasting technical success, and optimizing R&D portfolios.[1] However, systematic empirical evidence on how AI adoption changes firm-level innovative activity is limited. We address this gap by constructing a new firm-level measure of AI adoption and linking it to comprehensive patent data to study how AI affects the quantity, quality, and strategic composition of corporate innovation. Moving beyond innovation output, we examine how AI shapes novelty, technological direction, and the reach and scope of inventions.

Existing theoretical frameworks offer mixed predictions about how AI affects firm innovation. Some studies suggest that AI adoption can enhance corporate innovation by improving firms' ability to process complex and diverse information. For instance, Cockburn, Henderson and Stern (2018) argue that AI technologies support innovation by expanding firms' technological capabilities and accelerating knowledge discovery through automation of routine tasks and effective data analysis. Similarly, Mühlroth and Grottke (2020) present an AI-driven data mining framework that enables firms to identify emerging technology trends via automated query generation, data collection, topic modeling, and visualization. Empirically, Babina, Fedyk, He and Hodson (2024) find that firms investing in AI experience higher growth through increased product innovation.

In contrast, other studies express more cautious views regarding AI's impact on innovation.

[1] For instance, Microsoft employed Azure Quantum Elements, an AI-powered research platform, to accelerate scientific discovery in chemistry and materials science. To design a solid state battery electrolyte that outperforms those engineered decades ago, scientists face the challenge of screening a vast number of potential molecular combinations that can create an electrolyte (McPhee, 2024). AI models can help overcome this challenge by predicting the properties of electrolyte mixtures at scale. Scientists in Microsoft first used AI models to screen 32.6 million hypothetical materials and reduced the pool to about 500,000. Building on this initial screen, they coupled the models with physics based simulations, which narrowed the set to about 800 and then to about 150. In the final stage, AI property-prediction filters together with expert criteria selected 18 laboratory candidates, from which the lead electrolyte was synthesized and validated. This end-to-end discovery process, which would typically require years, was completed in less than nine months with the assistance of AI (McPhee, 2024; Microsoft, 2024). Microsoft subsequently filed a patent application covering the resulting electrolyte compositions (publication No. US2025/0105343 A1).

Brynjolfsson, Rock and Syverson (2019) and Azoulay, Krieger and Nagaraj (2025) argue that realizing the benefits of AI typically requires substantial complementary investments, such as cloud infrastructure, specialized human capital, and AI-oriented business processes, which entail large up-front and recurrent costs. Since these fixed and ongoing outlays can be particularly burdensome for smaller firms, AI adoption may divert limited resources away from other innovation investments, dampening aggregate innovative activity. Moreover, if AI's gains accrue disproportionately to a narrow set of early adopters, strong appropriability and scale effects can raise entry barriers and slow diffusion, discouraging innovation (Brynjolfsson et al., 2019; Cockburn et al., 2018). Additionally, AI's dependence on historical data to identify patterns and trends may reinforce existing innovation trajectories and limit the scope for exploratory innovation (Agrawal, McHale and Oettl, 2024). Furthermore, mass-produced, AI-generated disclosures can expand the body of citable prior art, rendering otherwise patent-worthy inventions unpatentable, reducing the expected probability of securing patent protection, and thereby weakening incentives to invest in invention (Yordy, 2021). Taken together, these concerns underscore that the overall effect of AI on firm innovation remains theoretically ambiguous and warrants empirical investigation.

One challenge in assessing AI's impact on innovation is measuring firm-level AI adoption. Prior studies often use AI-related job postings as a proxy for AI investment (Babina et al., 2024). This approach may understate true adoption if firms do not publicly advertise AI positions or already employ workers with relevant skills. Another common approach constructs task-level exposure to AI using Occupational Information Network (O∗NET) data and aggregates them to the occupation-, firm-, or industry-level (Eloundou, Manning, Mishkin and Rock, 2024; Acemoglu, 2025; Eisfeldt, Schubert and Zhang, 2025; Felten, Raj and Seamans, 2021). While informative, these proxies capture potential exposure rather than realized, firm-specific AI adoption. Leveraging newly available historical data on IT product installations, we construct a direct, usage-based measure of firm-level AI adoption that identifies adoption timing and enables analysis of its effects on firm-level innovation.

Given the staggered adoption of AI across firms, we apply a stacked difference-in-differences (DiD) approach that compares changes in innovation outcomes for cohorts adopting AI at different times (AI-adopting firms) and those not yet adopting (non-adopting firms).[2] Using data on granted patents filed by publicly listed U.S. firms to measure innovation outcomes, we find that AI-adopting firms experience increases in both the quantity and quality of patents compared to non-adopting firms. Patents produced by AI-adopting firms after adoption receive more citations than those of non-adopting firms, indicating higher scientific and technological value. Furthermore, AI adoption also influences firms' innovation strategies 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. Finally, 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 a greater number of claims, potentially reflecting a strategic effort to broaden the scope of their exclusive rights and enhance legal protection.

A natural concern is that AI adoption may be an endogenous choice and that unobservable factors could simultaneously influence both a firm's decision to adopt AI and its innovation performance. For instance, firms that invest heavily in innovation might be more likely to adopt AI and might have achieved better outcomes even without AI. To alleviate this concern, we first compare average R&D expenditure between AI-adopting and non-adopting firms and find that, on average, AI-using firms spend significantly less on R&D than their non-using counterparts. In addition, our stacked DiD estimations include firm-by-cohort and year-by-cohort fixed effects, which help isolate the effect of AI adoption from unobservable time-

[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.

invariant firm heterogeneity and cohort-specific shocks.

We further estimate a cohort-specific event study with event-time leads and lags around the adoption year to trace the dynamic effects of AI adoption. The results show that innovation increases following adoption and does not precede it. Moreover, a placebo test in which AI adoption is randomly assigned shows that these fictitious adopters do not outperform non-adopting firms in innovation in the absence of actual AI adoption, indicating that our results are not driven by selection or unobserved firm characteristics correlated with AI adoption. To further mitigate the self-selection concern that AI adopters differ from non-adopters in observed characteristics, we implement entropy balancing to ensure that AI-adopting and non-adopting firms are comparable in terms of key pre-treatment firm characteristics. This reweighting approach enhances the validity of our comparisons by addressing observable differences that may simultaneously influence both AI adoption and innovation outcomes.

To address remaining endogeneity concerns, we also employ a control function (CF) approach using firms' exposure to AI-strong universities developed by Babina et al. (2024) as an instrumental variable (IV) for AI adoption. This instrument satisfies the relevance condition, as firms with well-established hiring networks with AI-strong universities are more likely to access AI-trained talent and benefit from AI-related knowledge spillovers, thereby increasing their likelihood of adopting AI technology. It plausibly satisfies the exclusion restriction because the instrument is measured in 2010 before the commercial boom in AI and is not simply capturing institutional elite status. To further disentangle the effect of AI-specific talent from general human capital quality, we control for firms' exposure to top-10 universities following Babina et al. (2024). Our results on innovation quantity, quality, and strategic orientation remain robust after accounting for potential endogeneity in AI adoption.

We further explore the mechanisms through which AI influences corporate innovation. Our analysis indicates that AI promotes innovation through improved efficiency, changes in inventor composition, and enhanced recombination of knowledge and ideas. Interestingly, the cross-sectional analyses show that the impact of AI is generally more pronounced among firms

with a focused business scope. This pattern is contrary to our prior that AI would disproportionately benefit diversified firms operating in more complex information environments, given its capacity to process and learn from large volumes of data. One plausible explanation is that focused firms tend to generate high-quality, coherent, and consistent data from their operation, which enhances the effectiveness of AI applications. This finding is consistent with the view that the quality of firm-specific datasets plays a critical role in determining the performance of AI technologies (Brynjolfsson and McAfee, 2017; Brynjolfsson et al., 2019). Other factors that may enable focused firms to outperform include deeper domain expertise, less internal bargaining, and fewer coordination frictions.

Moreover, we find that firms facing intense product market competition tend to use AI to generate more original patents and a larger share of exploitative patents. This suggests that competitive pressure drives AI-adopting firms to recombine knowledge across diverse technological areas while simultaneously deepening expertise along their existing trajectories, allowing them to differentiate from competitors more efficiently. Additionally, the innovation effects of AI adoption are not stronger for larger or financially unconstrained firms, even though these firms are typically viewed as better positioned to bear the fixed costs of AI adoption, build the necessary data and computing infrastructure, and exploit AI's returns to scale.

We further check the robustness of our results via entropy balancing, placebo tests, and using never-adopting firms as an alternative control group. Our obtain similar results. We also show that our results are not primarily driven by firms operating in high-tech sectors and that AI's impacts on innovation are not confined to AI-related patents, underscoring its broader economic significance.

This paper contributes to the growing literature on the economic implications of artificial intelligence. While prior studies have examined AI's impact on firm growth and product innovation (Babina et al., 2024), technological trends (Mühlroth and Grottke, 2020; Webb, 2019; Agrawal, McHale and Oettl, 2019), patent examination (Zheng, 2025), firms' disclosure strategies (Cao, Jiang, Yang and Zhang, 2023), firm values (Eisfeldt et al., 2025), industry

entry and exit dynamics (Lu, Phillips and Yang, 2024), productivity and labor markets (Brynjolfsson, Li and Raymond, 2025; Chen and Wang, 2024; Jiang, Park, Xiao and Zhang, 2025), comprehensive evidence on how AI affects corporate innovation remains limited. Using a novel firm-level measure that directly captures AI adoption, our study extends this literature by showing that AI adoption systematically reshapes innovation, enhancing both quantity and quality and shifting its strategic direction. These effects are more pronounced for firms with a focused business scope, but not for larger firms with more resources or financially unconstrained firms.

The literature on the development and impact of AI technologies suggests that success hinges on data availability, computing capacity, and integration into core operations (Brynjolfsson and McAfee, 2017; Cockburn et al., 2018; Damioli, Van Roy and Vertesy, 2021). Recent survey studies underscore that high-quality data for building machine learning models are critical to the performance of AI (Zha, Bhat, Lai, Yang, Jiang, Zhong and Hu, 2025; Whang, Roh, Song and Lee, 2023). Consistent with this view, we provide new firm-level evidence that firms with a focused business scope realize larger innovation gains from AI adoption, as they tend to produce higher-quality, better-integrated, and more internally aligned data, accumulate deeper domain expertise, and face lower coordination and integration costs.

Our findings also speak to the ongoing debate on whether AI can profoundly alter the innovation production function and drive long-run economic, social, and technological changes (Cockburn et al., 2018; Alderucci, Branstetter, Hovy, Runge and Zolas, 2020). For instance, Cockburn et al. (2018) argue that advances in AI can enable a novel approach to conducting scientific and technical research, with economy-wide implications that diffuse across sectors. In contrast, Ouellette, Fang and Ouellette (2025) argue that the use of AI in patent drafting may lead to an overwhelming volume of low-quality patent applications, undermining examination efficiency and threatening the integrity of the patent system. Brynjolfsson and McAfee (2017) argue that the narrow scope and rivalrous nature of AI leads to an "AI gold rush". Meanwhile, Damioli et al. (2021) contend that AI remains in an early stage of economic maturity and it takes

more time before its productivity effects become clear. We add to this literature by showing that AI enhances firm innovation through greater efficiency, changes in inventor composition, and more effective knowledge recombination, suggesting a reshaping of the innovation process within firms.

2 Theoretical Background

The theoretical expectations about AI's influence on firm innovation are ambiguous. On the one hand, Mühlroth and Grottke (2020) suggest that AI supports firms in embedding strategic foresight into the innovation process, which helps detect emerging technology trends and innovations at an early stage. Ludwig and Mullainathan (2024) argue that AI can process large, complex, and unstructured data, such as satellite imagery, audio recordings, and online behavior, that were previously difficult for humans to analyze effectively, thereby uncovering subtle but meaningful patterns that humans may overlook. AI's ability to predict high-impact technological directions by analyzing vast datasets could stimulate innovation by identifying promising areas for research and development, leading to the creation of patents with greater technological value.

By drawing on various data sources, AI can help generate novel hypotheses and uncover relationships that foster more creative and original innovation (Haefner, Wincent, Parida and Gassmann, 2021). Beyond its advantage in data processing, AI may also transform the knowledge production function by allowing users to retrieve, represent, and combine knowledge from diverse technological domains in an automated manner (Agrawal et al., 2019). Further, AI may steer innovation towards the more efficient exploitation of existing knowledge bases. By integrating AI-driven insights into established technological expertise, firms can streamline experimentation, accelerate research and development, and reduce uncertainty in their innovation processes (Tranchero, 2024). Such exploitative innovation is particularly valuable in competitive environments, where efficiency is critical for sustaining an advantage.

In these settings, AI may allow firms to recombine current competencies and respond more rapidly to market and technological changes (Wamba-Taguimdje, Wamba, Kamdjoug and Wanko, 2020).

Consistent with AI's capability to process alternative data, firms can also leverage AI to efficiently search and analyze large volumes of text, which is particularly valuable for interpreting complex patent claim language and visualizing connections across intellectual property assets (Alderucci and Sicker, 2019; Setchi, Spasić, Morgan, Harrison and Corken, 2021). Since patent claims are essential for identifying prior art, assessing patentability, deterring potential infringement, and formulating legal strategies (Alderucci and Sicker, 2019), AI tools can assist firms in optimizing the design of their patent claims.

On the other hand, previous studies caution that AI adoption may introduce risks that might hinder innovation. The broad applicability of AI could lead to underinvestment in innovation due to coordination failures and externalities across sectors, ultimately slowing overall technological advancement (Bresnahan and Trajtenberg, 2017; Cockburn et al., 2018). Strong appropriability by early AI leaders may create bottlenecks, discouraging broader experimentation and innovation across sectors (Cockburn et al., 2018). Additionally, the complexity and significant resource requirements associated with AI deployment may divert managerial attention and firm resources from other critical innovative activities (Brynjolfsson et al., 2019), thereby reducing overall innovation capacity.

Furthermore, AI may diminish collective novelty by generating ideas that are less unique and more similar to one another (Doshi and Hauser, 2024). Agrawal et al. (2024) model AI-assisted discovery as prioritized search that ranks hypotheses by predicted success; without an explicit exploration objective, it tends to exploit existing data and neglect unexplored areas. AI's reliance on historical datasets for training may steer innovation toward established solutions and away from explorative directions. Consequently, heavy reliance on AI may encourage incremental innovation rather than fostering original or groundbreaking ideas. Moreover, AI-assisted patent drafting could undermine patent quality and threaten patent

system integrity by generating an overwhelming volume of low-quality applications that strain examination processes (Ouellette et al., 2025). In particular, automated drafting may lead to inadequate disclosures and improperly granted patents, with adverse consequences for the quality of future innovation. Taken together, these concerns suggest that AI adoption may not necessarily enhance, and could even constrain, firms' innovative performance.

3 Sample and Empirical Framework

3.1 Data and Sample

To measure firms' innovation activities, we collect patent, citation, and technology class data from the United States Patent and Trademark Office (USPTO).[3] Firm financial data are obtained from Compustat. We exclude firms in financial industries (SIC code 6000 to 6999) 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 Compustat firms, we construct a comprehensive linking table by integrating link tables from Kogan, Papanikolaou, Seru and Stoffman (2017), Wharton Research Data Services (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 retain only those matches for which we have high confidence.

The data on AI adoption come from the SWZD Historical Company Intelligence Database, which tracks historical site-level IT product purchases collected during IT consulting engagements. According to SWZD, this database captures approximately 90% of global IT purchasing

[3] We use the USPTO data updated on May 5, 2025.

activity, offering a comprehensive view of technology 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.

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 et al., 2001).

3.2 AI adoption measure

We use detailed site-level IT installation records for U.S. firms from the SWZD Historical Company Intelligence Database to measure firm's adoption of AI technologies. This database provides comprehensive information on IT products and their installation dates across various firm sites. The specificity and granularity of this data make it well-suited for tracking the timing and diffusion of AI technologies across firms. By identifying the earliest observed installation date of a machine learning product at any of the firm's sites, we measure the onset of AI adoption at the firm level. A firm is classified as an AI-adopting firm if it adopts any AI technology during the sample period; otherwise, it is categorized as a non-adopting firm. This usage-based approach enables us to track AI adoption through the actual deployment of technologies within firms, offering a more direct measure of AI adoption.

3.3 Innovation measure

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 (Leone and Reichstein, 2012; Lerner and Seru, 2022).

3.3.1 Quantity and quality

To capture the quantity of innovation outcomes, we count the number of patents filed by a firm in a given year. We proxy patent quality with the number of forward citations, which reflects a patent's technological impact (Hall, Jaffe and Trajtenberg, 2005). Given that the raw number of received citations is subject to truncation bias since citations take time to accrue, we follow Hall et al. (2001) and Lerner and Seru (2022) to adjust forward citation counts based on the patent's filing year and tech class. The metric is computed as follows:

$$Adjusted\ citation_p = \frac{Citation_p}{\sum_j^{N_{kt}} Citation_j / N_{kt}} \tag{1}$$

where $Citation_p$ is the number of citations received by patent p filed in year t, $Citation_j$ is the number of citations received by another patent j, filed in year t, within the same technology class k as patent p. N_{kt} denotes the total number of patents applied in year t within class k. This approach compares the number of citations a patent receives relative to other patents filed in the same year and technology class, thereby mitigating truncation bias.

3.3.2 Innovation strategy

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 index of backward citations cited by a patent:

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

where B_{pk} represents the ratio of backward citations made by patent p to patents in technology class k to the number of total backward citations made by patent p across all technology classes. 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 array of technological fields.

Distance measures the technology propinquity between a patent and the overall technology expertise of its owning firm (Akcigit, Celik and Greenwood, 2016a). Firstly, we compute the distance between technology classes, defined by the first two digits of the Cooperative Patent Classification (CPC) code. Specifically, the distance between two technology classes X and Y is based on the citations made to classes X and Y. We use $\#(X \cap Y)$ to denote the number of patents that cite at least one patent from both class X and Y (intersection) and use $\#(X \cup Y)$ to denote the number of patents that make citations to at least one patent from either class X or Y (union). Hence, the distance metric between class X and Y can be calculated as follows:

$$d(X, Y) \equiv 1 - \frac{\#(X \cap Y)}{\#(X \cup Y)} \tag{3}$$

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.[4]

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

$$d_t^{\iota}(p, i) \equiv \left[\frac{1}{\|\mathcal{P}_{it}\|} \sum_{p' \in \mathcal{P}_{it}} d(X_p, Y_{p'})^{\iota} \right]^{\frac{1}{\iota}} \tag{4}$$

[4]For robustness, we compute two versions of this distance metric for a given class pair in a given year: one based on citations in that specific year, and another based on cumulative citations up to that year. The two methods yield similar results. We report the results based on technology distances calculated from citations made up to the focal year.

where $\mathcal{P}_{it}$ denotes the set of all patents filed by firm i prior to the focal 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 focal 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 et al. (2016a) 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 as the set of (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 the filing date of the focal patent to the filing date of the focal patent. 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)..[5] This measure captures the extent to which a firm's innovation improves along the existing technology trajectory (Benner and Tushman, 2002)

On the contrary, explorative patents refer to inventions that are mainly built outside of a firm's existing knowledge base. A patent is classified as explorative if 80% and more of its citations are outside a firm's existing knowledge base (that is, neither belong to the firm's own patent filings nor were cited by the firm from the beginning of the fifth calendar year before the filing date of a focal patent to the filing date of the focal patent).[6] This measure gauges the extent to which a firm's innovation shifts to a new technology trajectory (Benner and Tushman, 2002).

[5] As a robustness check, we also use a 60% threshold and our results hold under the alternative cutoff.
[6] As a robustness check, we also use a 60% threshold and our results hold under the alternative cutoff.

3.3.2.3 Innovation reach: generality and patent claim

Generality reflects how widely a patent's impact spreads across different technology classes (Hall et al., 2001). Following Trajtenberg et al. (1997), we calculate generality as one minus the Herfindahl index of forward citations citing a patent:

$$Generality_p = 1 - \sum_{k}^{n_p} F_{pk}^2 \tag{5}$$

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). We follow Akcigit, Baslandze and Stantcheva (2016b) and Aghion, Akcigit, Bergeaud, Blundell and Hémous (2019) to count the number of claims contained in a patent application.

4 Empirical analyses

4.1 Summary statistics

We report innovation activities and firm characteristics for all firms in the sample in Columns (1)-(3) of Table 1. An average firm in our sample produces 24.685 patents and receives 0.521 forward citations per patent, adjusted by technology class and filing year. The average originality is 0.271, and the average technology distance is 0.242. On average, firms generate 49.90% exploitative patents and 14% explorative patents. The average generality is 0.093, and the average number of claims per patent is 12.262.

In our sample, 518 firms adopt AI technologies, while 2,082 firms never do so.[7] Columns (4)-(9) of Table 1 report innovation activities and firm characteristics of firms that adopt AI technology and firms that never adopt AI technology. 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. On average, an AI-adopting firm produces 65.971 patents, while a never-adopting firm produces 9.308 patents. The average adjusted citation count for patents filed by AI-using firms is 0.672, compared to 0.465 for patents filed by never-using firms. Patents from AI-using firms also display greater average originality (0.291 vs. 0.264) and technology distance (0.295 vs. 0.222). We further document that 64.30% (44.60%) of patents in AI-adopting (never-adopting) firms are exploitative, and 10.30% (15.30%) of patents in AI-adopting (never-adopting) firms are explorative. In addition, the average generality for patents from AI-adopting firms is 0.109, compared to 0.087 for patents from never-adopting firms. Patents from AI-using firms also contain more claims than those from never-using firms (14.079 vs. 11.585).

It is worth mentioning that AI-using firms, on average, spend less on R&D expenditure than never-using firms. The difference is both economically and statistically significant. This helps alleviate the potential concern that our findings would be biased if AI-adopting firms systematically spent more in R&D investments. We find that AI-adopting firms are larger, have higher Tobin's Q values, and hold less cash compared to never-adopting firms.

4.2 Empirical analyses

4.2.1 Baseline model

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

[7] Based on the Fama-French 49 industry classification, Appendix Table A.2 shows that AI-adopting firms are concentrated in industries such as computer software, electronic equipment, retail, and business services, and are least represented in non-metallic and industrial metal mining, agriculture, textiles, coal, fabricated products, and precious metals.

or over time (Baker, Larcker and Wang, 2022). To mitigate this bias, we follow Gormley and Matsa (2011, 2016) to use a stacked difference-in-differences approach to study the impacts of AI adoption on corporate innovation. In this setup, we define cohorts based on the year of AI adoption—referred to as the event year. For each cohort, we construct an event window spanning five years before to five years after the adoption year. Treated firms are those that adopt AI in a given 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 final stacked sample.

To empirically test how AI adoption affects innovation activities, we estimate the following model on the stacked sample, separately for each $h \in \{1, 2, 3\}$:

$$\mathbb{E}\left[Y_{i,j,c,t+h} \mid \cdot \right] = \exp\Big(\beta_h \, AI_{i,j,c,t} \;+\; \alpha_{i,c} \;+\; \psi_{t+h,c}\Big). \tag{6}$$

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. $\alpha_{i,c}$ and $\psi_{t+h,c}$ are firm-by-cohort and year-by-cohort fixed effects. 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 (6) by Pseudo-Poisson Maximum Likelihood (PPML), following Cohn, Liu and Wardlaw (2022). Our parameters of interest are β_h, which capture differences in changes in patent quantity, quality, and innovation strategy around AI adoption between adopting firms and the control groups at each horizon h.

4.2.2 Baseline results

4.2.2.1 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 $AI_{i,j,t,c}$, 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.

We further 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—filed within three years following a firm's adoption of AI technology. As shown in Panel B of Table 2, the coefficients on $AI_{i,j,t,c}$ are all positive and significant, which indicates that patents from AI-using firms receive more adjusted citations after AI adoption than those from non-adopting firms. The result suggests that AI-adopting firms develop patents of higher scientific value after using AI than non-adopting firms.

4.2.2.2 Innovation strategy

We further study whether the use of AI affects the innovation strategy of AI-adopting firms. We compare the innovation novelty (measured by originality and distance), innovation direction (measured by exploitative and explorative patents), and innovation reach (measured by generality and number of claims) between patents invented by AI-adopting firms before and after using AI with those produced by non-adopting firms.

Table 3 Panel A reports the results for the impacts of AI on innovation novelty. In Columns (1)-(3), we document that AI-using firms experience a notable increase in the average originality of their patents following adoption, compared to non-adopting firms. This effect is both statistically and economically significant. After three years of adopting AI technology, AI-using firms experience 21.53% ($(e^{0.195} - 1) \times 100\%$) increase in originality, compared to non-using firms. In Columns (4)-(6), we show that firms experience significant increases

in distance after AI adoption. After three years of adopting AI technology, AI-using firms experience 13.31% ($(e^{0.125} - 1) \times 100\%$) increase in distance, compared to non-using firms. The results in Panel A of Table 3 imply that after the adoption of AI technology, firms tend to produce patents that are built on a wider array of technology fields and that are more technologically distant to their previous patent portfolios, compared to non-adopting firms.

Panel B of Table 3 reports the results of how the use of AI is related to innovation direction. We present the results for exploitative and explorative patents in Columns (1)-(3) and Columns (4)-(6), respectively. Post AI adoption, firms generate more exploitative and explorative patents than non-adopting firms. In the third year after firms start to use AI technology, exploitative and explorative patents grow by 16.65% ($(e^{0.154} - 1) \times 100\%$) and 21.04% ($(e^{0.191} - 1) \times 100\%$), respectively. These results suggest that the use of AI helps adopting firms not only deepen the utilisation of their existing technological base, but also broaden it by exploring new technological knowledge.

Table 3 Panel C reports the results of how AI affects innovation reach. In Columns (1)-(3), we show that after adopting AI technology, patents produced by AI-adopting firms illustrate significantly higher generality than patents produced by non-adopting firms. In Columns (4)-(6), we also observe that patents of AI-using firms have a higher number of claims after AI adoption, compared to patents of non-using firms. Three years post adoption, patents of AI-using firms have approximately 17.48% ($(e^{0.161} - 1) \times 100\%$) more claims than patents of non-using firms. These findings indicate that firms adopting AI technologies produce patents with extended technological and legal scope, thereby shaping future innovation across a wider spectrum of technological fields.

4.2.2.3 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 with improving pre-adoption innovation are more likely to adopt AI, and 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 et al. (2022), let $t_{i,c}$ denote firm i's adoption year in cohort c and define the relative time $\ell \equiv t - t_{i,c}$. 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:

$$\mathbb{E}\big[\, Y_{i,j,c,t+1} \mid \cdot \,\big] = \exp\left(\sum_{\ell=-5}^{-2} \beta_\ell \, AI^{\ell}_{i,j,c,t} \;+\; \sum_{\ell=0}^{5} \beta_\ell \, AI^{\ell}_{i,j,c,t} \;+\; \alpha_{i,c} \;+\; \psi_{t+1,c} \right), \tag{7}$$

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 from 3 years before the treatment to 4 years after treatment. Panel A shows that the difference in patent quantity and quality between AI-adopting and non-adopting firms increases significantly two years after AI adoption. Panels B, C, and D report differences in patent strategy in terms of novelty, direction, and reach, respectively. In Panel B, we document that differences in patent originality and distance between AI-adopting and non-adopting firms become significant two years after adoption. Panel C suggests that AI-using firms produce notably more exploitative patents after two years' of AI adoption compared to non-using firms, while there are no significant changes in explorative patents. Panel D indicates that differences in patent generality become significant three years after AI adoption and the number of claims become significant one year

after AI adoption. Furthermore, we find no significant pre-adoption differences in innovation activity between AI adopters and non-adopters, mitigating concerns that the results are driven by reverse causality or pre-existing trends.

Our results also 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)).

4.2.3 Endogeneity

The baseline estimates may be biased by 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 absent AI adoption. In this case, these unobserved determinants jointly affect both adoption and innovation, confounding the estimated effect of AI adoption on innovation. To alleviate this concern, we implement a control function regression approach using an instrumental variable to isolate exogenous variation in AI adoption.

A control function regression approach is a useful method to address endogeneity, especially in non-linear models (Wooldridge, 2015). In the first stage, it uses one or more instrumental variable to isolate exogenous variation in the endogenous variable, providing separate variation in the generalized residuals. In the second stage, these generalized residuals are included as a control function, so that the endogenous variable is conditionally exogenous.

Compared with two-stage least squares (2SLS), the control function approach provides a more parsimonious way to handle non-linear models with endogenous explanatory variables (Wooldridge, 2015) and therefore is more suitable for our setting.

We follow Babina et al. (2024) to use *AI exposure* as an instrumental variable, which calculates the focal firm's exposure to AI-strong universities via the firm-university STEM worker hiring networks as of 2010. As noted by Fujii and Managi (2018), universities are key AI technology inventors in both the U.S. and China. Babina, Fedyk, He and Hodson (2023) find that firms that have higher initial shares of highly-educated and STEM employees are more likely to invest in AI. Moreover, Babina et al. (2024) document that the scarcity of well-trained AI talent is a major constraint on firms' ability to adopt AI technology. Therefore, our instrument satisfies the relevance condition: firms that maintain well-developed hiring pipelines with AI-strong universities in 2010 are more likely to attract and access high-quality AI talent, and consequently adopt AI technologies. We thus expect a positive relationship between a firm's ex-ante exposure to AI-strong universities and its likelihood of adopting AI technology.

The exclusion restriction may be challenged if firms' exposure to AI-strong universities in 2010 influences firms' innovative activities in ways other than AI adoption. However, we argue that this is less of a concern in our setting for several reasons. First, *AI exposure* is an ex-ante measure built on firms' hiring connections to these universities as of 2010 when the commercial use of AI had not yet become widespread (Babina et al., 2024). In addition, Babina et al. (2024) document that from 2005 to 2010, the proportion of graduates hired from AI-strong universities remains stable among firms connected to those institutions. This helps alleviate the endogenous concern that it is the anticipation of future needs for AI talent that encourages the hiring of AI-trained graduates in 2010. Second, while the large-scale commercial deployment of AI has only surged in recent years, academic research on artificial intelligence can be traced back to the 20th century. Given that AI-strong universities are identified using academic publications prior to 2010, it is relatively unlikely that it is the demand of AI-adopting firms that reversely

promotes AI research in universities prior to 2010. Third, access to AI talent and access to high-quality human capital do not necessarily overlap: of the 44 AI-strong universities, 11 (25) are in the U.S. News & World Report top 20 (top 50). To further address the possibility that access to top-tier human capital confounds the results, we include a control variable, *Top 10 exposure*, which measures the focal firm's exposure to top 10 universities via the firm-university STEM worker hiring networks as of 2010 (Babina et al., 2024).

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

$$\begin{aligned} AI^{*}_{i,j,c,t} &= \beta\, AI\ exposure_{i,j,c} \;+\; \theta\, Top\ 10\, exposure_{i,j,c} \;+\; \gamma_{j,c} \;+\; \tau_{t,c} \;+\; u_{i,j,c,t}, \\ AI_{i,j,c,t} &= \mathbf{1}\{AI^{*}_{i,j,c,t} > 0\}, \qquad \Pr\big(AI_{i,j,c,t} = 1 \mid \cdot\big) = \Phi\big(z'_{i,j,c,t}\boldsymbol{\delta}\big), \end{aligned} \tag{8}$$

where $AI\ exposure_{i,j,c}$ is exposure to AI-strong universities via firm–university STEM hiring networks as of 2010 (Babina et al., 2024); $Top\ 10\, exposure_{i,j,c}$ is exposure to the top-10 universities (2010 U.S. News & World Report) via the same networks. $\gamma_{j,c}$ and $\tau_{t,c}$ are industry-by-cohort and year-by-cohort fixed effects. We do not include firm fixed effects because the instrumental variable varies cross-sectionally and would be absorbed by them. Let $z_{i,j,c,t}$ stack the regressors and fixed effects, and let $\boldsymbol{\delta}$ denote the corresponding coefficient vector.

We compute the generalized residuals from the probit model as:

$$\widehat{Residual}_{i,j,c,t} = AI_{i,j,c,t}\,\frac{\phi\big(z'_{i,j,c,t}\hat{\boldsymbol{\delta}}\big)}{\Phi\big(z'_{i,j,c,t}\hat{\boldsymbol{\delta}}\big)} \;-\; \big(1 - AI_{i,j,c,t}\big)\,\frac{\phi\big(z'_{i,j,c,t}\hat{\boldsymbol{\delta}}\big)}{1 - \Phi\big(z'_{i,j,c,t}\hat{\boldsymbol{\delta}}\big)}, \tag{9}$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ denote the standard normal PDF and CDF, and $\hat{\boldsymbol{\delta}}$ is the first-stage estimate.

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

$$\mathbb{E}\big[\, Y_{i,j,c,t+h} \mid \cdot \,\big] = \exp\Big(\beta\, AI_{i,j,c,t} + \theta\, Top\ 10\, exposure_{i,j,c} + \kappa\, \widehat{Residual}_{i,j,c,t} + \gamma_{j,c} + \psi_{t+h,c}\Big). \tag{10}$$

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 innovation quantity, quality, strategy. For brevity, we report the estimates in year $t+3$ since the results for other years are similar. 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. Columns (2)-(9) present the second-stage regression results. We find that the positive effects of AI adoption persist across various innovation outcomes, including *patent count*, *adjusted citations*, *originality*, *distance*, *exploitative*, *generality*, and *claim*. Overall, these results suggest that differences in innovation quantity, quality, and strategy between AI-adopting and non-adopting firms remain robust after accounting for the endogeneity of AI adoption.

5 Channels

The results thus far suggest that AI adoption significantly affects adopters' innovation outcomes. In this section, we examine the potential channels through which AI technology influences innovation.

5.1 Efficiency channel

5.1.1 Inventor productivity

We begin by examining an efficiency channel through which AI can raise corporate innovation by increasing patent output generated from a given set of inventive inputs. We first focus on inventor productivity and discuss three related mechanisms through which AI can increase patent output per inventor. First, AI can automate routine, time-intensive tasks and process technical and alternative data at scale. Setchi et al. (2021) show that AI can reduce the time

and cost of prior-art searches by efficiently screening large volumes of patent documents. Second, AI can lower the costs of accessing and internalizing relevant knowledge within the firm by codifying and disseminating effective practices, which may particularly benefit less-experienced inventors. Brynjolfsson et al. (2025) find that a generative AI-based conversational assistant increases the productivity of novice and lower-skilled customer-support agents by summarizing top performers' practices and diffusing them across workers. Third, AI can accelerate knowledge search and recombination by helping inventors retrieve, synthesize, and connect disparate information, thereby increasing the rate at which they generate patentable ideas. For instance, AI can help identify emerging technological trends (Mühlroth and Grottke, 2020) and support hypothesis generation (Ludwig and Mullainathan, 2024).

To explore this channel empirically, we define *inventor productivity* as the average number of granted patents filed by inventors of firm i in year t. We first calculate the sum of all patents filed by the focal inventor at the inventor-year-level. For patents with multiple inventors, we assign each inventor an equal fraction of the patent (Moretti, 2021). We then aggregate the inventor–year patent counts to the firm–year level.[8] Using Equation 6, we compare inventor productivity in AI-adopting firms before and after AI adoption with that of non-adopting firms. Columns (1)-(3) in Panel A of Table 5 reports the results. We find that inventor productivity in AI-adopting firms increases significantly following AI adoption, relative to that in non-AI firms. Compared with non-using firms, AI-using firms experience a 14.91% ($(e^{0.139} - 1) \times 100\%$) increase in inventor productivity three years after adoption. Our results suggest that AI promotes corporate innovation through the mechanism of inventor productivity.

5.1.2 Innovation efficiency

AI may also affect corporate innovation through innovation efficiency, the effectiveness with which firms convert R&D spending into patents. AI can enhance innovation efficiency by improving the screening of research opportunities and increasing the expected discoveries

[8] As a robustness check, we compute *inventor productivity* as the ratio of the number of patents filed by firm i in year t to the number of inventors who file patents in firm i in year t. Our results remain similar.

per experiment. For example, using AI to predict the bioactivity of candidate molecules substantially improves the productivity of early-stage drug screening (Cockburn et al., 2018). Merchant, Batzner, Schoenholz, Aykol, Cheon and Cubuk (2023) argue that AI can accelerate material discovery by allowing out-of-distribution exploration in in previously unseen chemical spaces and by improving predictive performance on downstream tasks. Bennett, Watson, Ragotte, Borst, See, Weidle, Biswas, Yu, Shrock, Ault et al. (2025) claim that AI models can substantially improve antibody discovery and development by enabling faster and more cost-effective identification of promising candidates compared with approaches such as animal immunization and random library screening.[9]

By reducing information frictions in project selection and enabling better allocation of R&D resources, AI can help firms direct investment toward higher-potential projects and generate more patent output per dollar spent. Yoo, Jung and Jun (2023) show that using machine learning models to predict the performance of R&D projects improves the efficiency of funding decisions. Efficiency gains from better project selection and fewer avoidable failure modes can lead to higher patent quality. Finally, AI may enable more efficient R&D experimentation. By automating screening, candidate generation, and experiment design, AI lowers the marginal cost of search and expands the feasible search space for experimentation, raising innovation efficiency (Cockburn et al., 2018).

To investigate this channel, we follow Gao and Chou (2015) to define *innovation efficiency* as the firm-year ratio of the number of patents to R&D expenditure. We use Equation 6 to compare the innovation efficiency in AI-using firms before and after using AI technology with that of non-using firms. Columns (4)-(6) in Panel A of Table 5 shows that AI-using firms experience a significant increase in innovation efficiency relative to non-adopting firms. Compared with non-adopting firms, innovation efficiency in AI-adopting firms rises by 18.07% ($(e^{0.166}-1)\times 100\%$) three years after adoption. Our results indicate that AI promotes

[9] McKinsey & Company (2023) suggests that AI models can accelerate research operations by both speeding up each simulation and enabling multiple tests simultaneously. Based on a survey of biopharma R&D executives, Accenture (2023) finds that AI-led discovery strategies can help firms shorten discovery cycle times by up to two-thirds.

corporate innovation through a higher rate of conversion from inventive effort to patentable outcomes.

5.2 Inventor composition channel

AI adoption may affect not only inventor productivity, but also the firm's optimal choice of inventive labor inputs along both the extensive margin (scale) and the intensive margin (composition of skills and experience). An important insight from the technology and labor literature is that information technology complements workplace reorganization and skilled labor, so adopting new technologies is often accompanied by changes in workforce composition (Bresnahan, Brynjolfsson and Hitt, 2002). The direction of these adjustments depends on whether a technology primarily replaces human task performance (substitution) or raises the productivity of humans performing those tasks (complementarity).

This perspective implies that AI can reshape the innovation process by reallocating tasks within R&D and, in turn, shifting demand across inventive skills. In the task-based framework of technological change, computerization substitutes for routine, codifiable tasks while complementing non-routine problem solving and complex communication (Autor, Levy and Murnane, 2003). To the extent that AI expands the set of activities that can be automated or reliably augmented, firms have incentives to reorganize inventive work by shifting routine components from labor to capital and reallocating human effort toward tasks where researchers retain comparative advantage. For example, if AI automates or substantially reduces the time cost of routine R&D inputs, such as documentation, code templating, prior-art and literature search, and basic experimentation, then inventors can be redeployed towards higher-value activities, including framing research questions, generating and testing hypotheses, debugging complex designs, and coordinating across technical domains.

These adjustments could lead to changes in the composition of the inventor workforce. First, by lowering fixed costs of producing and codifying inventions, AI may increase the size of the inventor pool and encourage entry of first-time inventors. Second, if AI raises

the productivity of non-routine analytic and evaluative tasks, the marginal product of skilled inventive labor increases, leading firms to tilt team composition toward inventors who can effectively deploy, supervise, and validate AI-augmented workflows. Third, if AI automates routine components of invention while complementing evaluation and integration, the internal division of labor in R&D may shift, changing field specialization across inventors and teams.

To explore this channel, we construct four variables to measure inventor composition. We define *inventor pool* as the total number of unique inventors who file at least one patent with firm i in year t. 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. *AI-experienced inventor* for firm i in year t is measured as 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 are AI-experienced. An inventor is classified as AI-experienced if she has filed at least one patent 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 up to year t-1. *Specialization* computes the average specialization scores of all inventors who file at least one patent with firm i in year t. An inventor's specialization score in year t is the Herfindahl 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 class.

Using these measures as dependent variables, we estimate Equation 6 to compare changes in inventor composition in AI-adopting firms before versus after adoption relative to non-adopting firms. The estimations include year–cohort fixed effects to control for aggregate trends in inventors and common calendar-year shocks, such as overall trends and macroeconomic conditions, within each cohort-specific event-time window, as well as firm–cohort

fixed effects to control for unobserved, time-invariant firm characteristics, such as persistent management quality, corporate culture, or technology level.

Panel B of Table 5 shows that AI-using firms experience significant increases in *inventor pool*, *first-time inventor*, *AI-experienced inventor*, and *specialization* compared with non-using firms. These results suggest that AI influences corporate innovation through changes in inventor composition. Specifically, in comparison with non-adopters, adopters employ more inventors after AI adoption, including inventors with no prior patenting experience as well as those with prior experience in AI patents. Moreover, AI adoption is associated with a relative increase in inventor technological specialization in AI-adopting firms compared with non-adopting firms, aligning with evidence that AI adopters tilt innovation toward exploitation of the firm's existing technological strengths.

5.3 Knowledge recombination channel

Another mechanism emphasized in the innovation literature is that invention is often recombinant: new patents frequently arise from reconfiguring existing ideas or technological components (Weitzman, 1998; Fleming, 2001). In this view, the production of patents depends on access to relevant prior knowledge and the ability to search over and evaluate novel combinations under technological uncertainty (Fleming, 2001). As the knowledge frontier expands, inventive activity increasingly requires coordination among specialized contributors, contributing to the long-run shift toward team-based patenting (Jones, 2009; Wuchty, Jones and Uzzi, 2007).

If AI reduces the cost of searching, evaluating, and integrating interdependent technological components, it should facilitate recombinant invention by expanding the set of knowledge combinations that can be feasibly explored and validated. As the number of relevant components grows, the combinatorial search space expands, increasing the marginal cost of generating, screening, and validating candidate solutions. AI can relax this constraint by accelerating recombination search, including the generation and prioritization of candidate

combinations, and improving screening and validation under uncertainty.[10]

This recombination channel has implications for inventive labor inputs. By expanding the scope of feasible combinations, AI can increase the marginal value of assembling complementary expertise within a project by enabling teams to cover more distinct knowledge domains. At the same time, by lowering cross-domain coordination and knowledge-transfer frictions, AI can raise the returns to maintaining broader technical breadth within teams. Taken together, these forces predict changes in team size and patent knowledge inputs following AI adoption.

We measure knowledge recombination from four perspectives. We define *team size* as the firm-year average number of inventors per patent, and *unique inventor team* as the number of distinct inventor teams, defined by unique inventor combinations, that file at least one patent with firm i in year t. For firm i in year t, *core skill breadth* is defined as 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. We follow (Li and Wang, 2023) to define an inventor's core skill in year t 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*. *New cite ratio* for firm i in year t is calculated as the is the average of the patent-level new cite ratio across the firm's patents filed in year t. The new cite ratio for patent p filed by firm i in year t is defined as the share of backward citations made to new knowledge, that is, knowledge outside the inventor team's existing knowledge, relative to the total number of backward citations to U.S. patents, foreign patents, and other references (Fitzgerald and Liu, 2020). The inventor team's new knowledge consists of citations to U.S. patents, foreign patents, or other references that have neither been previously cited nor produced by any member of the inventor team prior to the filing date of patent p. The inventor team's existing knowledge comprises all U.S. patents, foreign patents, and other references that

[10]For example, Exscientia reports using an AI-enabled discovery platform within a design–make–test–learn workflow in the discovery of clinical candidate EXS21546. A related patent disclosure includes bioactivity data for 46 exemplified molecules, which Exscientia reports as a subset of 163 molecules synthesized and assayed during discovery, illustrating how AI can help generate and prioritize candidate recombinations for iterative experimental validation (Wills, 2022).

have been previously cited or produced by at least one team member before the filing date of patent p.

Using these measures as dependent variables, we estimate Equation 6 to compare changes in knowledge recombination in AI-adopting firms before and after AI adoption relative to non-adopting firms. Panel C of Table 5 indicates that, relative to non-adopting firms, AI-adopting firms show significant increases in *team size*, *unique inventor team*, *core skill breadth*, and *new cite ratio*. These results indicate that AI affects corporate innovation through the channel of knowledge recombination. Following AI adoption, inventors in AI-using firms form larger and a greater number of unique inventor teams and cite knowledge that is new to all team members, compared with inventors in non-using firms. In addition, inventor teams in AI-adopting firms exhibit broader core technical skill coverage after adoption relative to non-adopting firms, consistent with greater cross-domain capability and higher capacity to integrate knowledge inputs within the team.

6 Cross-Sectional Heterogeneity

6.1 Business focus

A growing literature emphasizes that AI and other data-intensive technologies are particularly valuable in data-rich and complex information environments. Theoretical models that view AI as an improvement in prediction show that the gains from better prediction are larger when decisions are made in complex environments with many potential states and abundant data (Agrawal, Gans and Goldfarb, 2018). Mihet and Philippon (2019) argue that big data and AI technologies can be viewed as intangible assets, search and matching technologies, and forecasting technologies, in ways that strengthen the advantages of data-rich firms. Consistent with this view, Begenau, Farboodi and Veldkamp (2018) show that big-data technologies disproportionately benefit firms with more economic activity and longer histories, because they generate more data that can be processed and exploited. Accordingly, we conjecture that

AI adoption benefits diversified firms operating in richer data and more complex information environments.

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 and no additional segments 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 6. The estimated coefficients on this interaction term capture the differential impacts of AI on corporate innovation for firms with different business focus. 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 strategy in the third year following AI adoption. Column (1) presents estimates for innovation quantity, Column (2) for innovation quality, and Columns (3) through (6) for various innovation strategy measures. 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 richer and more complex data by enabling them to process and exploit large volumes of complex information.

Several factors may help explain why firms with a more focused business scope are better positioned to utilise AI to generate innovation than diversified firms. First, focused firms are more likely to generate higher-quality and more coherent internal data. Because data quality is central to AI performance, such high-quality and internally consistent data enhance AI effectiveness by reducing noise and inconsistencies in model training and fine-tuning (Zha

et al., 2025; Whang et al., 2023; Brynjolfsson and McAfee, 2017). Second, the productivity of AI depends on complementary assets such as data architecture and domain expertise (Brynjolfsson et al., 2019; Azoulay et al., 2025). Business focus deepens domain expertise and tacit knowledge, improving problem formulation, labeling, and feature engineering, thereby increasing the productivity of AI in invention. Third, AI adoption requires complementary fixed investments in data infrastructure. With fewer competing units, focused firms face less internal bargaining and capital-allocation distortion, enabling more targeted and scalable deployment of these complements (Rajan, Servaes and Zingales, 2000). Fourth, effective AI implementation requires coordination across engineering, product, legal, and compliance. Focused firms face fewer coordination frictions and lower integration costs, allowing faster and more complete embedding of AI in the invention process. Overall, these mechanisms suggest that focused firms can capture greater innovative returns from AI by combining higher-quality data, more targeted complementary investments, and smoother organisational integration.

6.2 Product market competition

Product market competition may also shape the impact of AI on corporate innovation. Under high product market competition, firms face strong pressure to act swiftly and differentiate themselves from their competitors (Arts, Cassiman and Hou, 2025). In such environments, AI can play a pivotal role by enabling firms to process data more effectively, identify emerging technological opportunities, and recombine knowledge more effectively, thereby generating more differentiated innovation outcomes. Furthermore, the pressure of product market competition might lead firms to favor exploitative innovation, which builds on existing technological capabilities and yields incremental improvements, over explorative innovation, which involves greater uncertainty and typically requires longer time horizons. As noted by Jansen, Van Den Bosch and Volberda (2006), exploitative innovation is more conducive to short-term financial performance in highly competitive settings, making it a strategic priority for firms under intense market pressure. By leveraging familiar technologies and established

knowledge bases, firms can generate innovative outputs while reducing the risks associated with pursuing untested innovation paths. Based on this logic, we expect that AI-using firms in highly competitive product markets are more likely to leverage AI to generate novel innovation while still building on their existing technological strengths.

We follow Hoberg and Phillips (2016) to define *high competition*, which equals 1 if a firm's TNIC-based Herfindahl-Hirschman Index (HHI) is equal to or below the 2015 average, and 0 otherwise. To assess the impact of product market competition on the relationship between AI adoption and innovation, we introduce an interaction term between *AI* and *High competition* in Equation 6. We include firm-cohort-high competition and year-cohort-high competition fixed effects to account for unobserved heterogeneity across firms, market structures, and time.

Panel B of Table 6 reports the results of the differential impacts of AI on corporate innovation for firms facing different product market competition. The coefficients on the interaction term indicate that competitive pressure amplifies the impact of AI adoption on *originality* and *exploitative*, but does not significantly alter AI's influence on other innovation outcomes. In other words, firms leverage AI to improve existing strengths and generate novel knowledge combinations under high competition.

6.3 Firm size

A natural concern is that the estimated effect of AI adoption on corporate innovation is confounded by firm size. Larger firms systematically differ from smaller ones, both in their propensity to adopt AI and in the benefits they can extract from it, for reasons unrelated to AI itself. 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 adoption (Babina et al., 2024; Lu et al., 2024). At the same time, large firms may inherently generate more innovation regardless of AI, owing to structural advantages such as economies of scale, greater market influence, and broader patent portfolios. If this is the case, our baseline estimates could simply capture the behavior of large firms.

To examine whether our results are driven by firm size, we augment Equation 6 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. Similarly, we 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 C of Table 6 reports the coefficients on this interaction, which capture the differential impact of AI on corporate innovation for large versus smaller firms. The interaction term is statistically insignificant for all innovation outcomes except *adjusted citation*. These results indicate that AI's effects on innovation quantity and innovation strategy do not differ systematically across firm sizes. For innovation quality, as captured by *adjusted citation*, the interaction coefficient is only weakly significant, indicating at most weak evidence of modestly stronger gains among larger adopters. Overall, these results suggest that our main findings are not primarily driven by big firms.

6.4 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 financing frictions explain weaker innovation among non adopters compared with AI adopters, we introduce an interaction term between *AI* and *Constrained firm* in Equation 6. Using the financial constraint measure developed by Hoberg and Maksimovic (2015), we define *Constrained firm* as an indicator variable equal to 1 if a firm's risk of

delaying investments because of liquidity issues is at least as large as the 2015 industry average, and 0 otherwise. In addition, we control for firm-cohort-constrained firm and year-cohort-constrained firm fixed effects to absorb differences related to firms' financing constraints across cohorts and over time. In Table 6 Panel D, we present estimates for the interaction term that captures the differential impact of AI on corporate innovation for firms facing different financial constraints. Our results show that the coefficients on *AI*×*Constrained firm* are significant for *adjusted citation* and *distance*, indicating that financially constrained firms use AI to generate patents that are both higher in quality and more technologically distant from their existing patent portfolios. The estimated coefficients on *AI*×*Constrained firm* are insignificant for other innovation measures. These results are inconsistent with financial constraints being the driver of weaker innovation among non adopters compared with AI adopters.

6.5 Technological capabilities

Firms differ in their technological capabilities. In particular, high-tech firms with more developed R&D infrastructure and inventor expertise, might be more likely to adopt AI and inherently more innovative. Consistent with this, Webb, Short, Bloom and Lerner (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 rather than the effect of AI adoption itself.

To alleviate the concern that AI-adopting and non-adopting firms differ systematically in their underlying technological capabilities, we augment Equation 6 with an interaction term between *AI* and *High-tech firm* to examine whether the effect of AI adoption on innovation differs 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") (Babina et al., 2024). We further include firm-cohort-high-tech firm and year-cohort-high-tech firm fixed effects 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 E of Table 6 reports that the coefficients on *AI*×*High-tech firm* are mostly insignificant, indicating that we do not find evidence that high-tech firms 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 primarily by pre-existing differences in technological capabilities between AI-adopting and non-adopting firms, as opposed to the effect of AI itself.

7 Robustness checks

7.1 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-using) and control (non-using) 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 reweighted (entropy balanced) sample, we re-estimate the baseline stacked DiD model (Equation 6) and report the results in Panel A of Table 7. 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.

7.2 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 6 using this placebo sample.

Panel B of Table 7 reports the estimation results. Across all specifications, the coefficients on AI are significantly negative or statistically insignificant. These findings suggest that our results are unlikely to be driven by pre-existing differences or other confounding effects.

7.3 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 6) using only never-treated firms as the control group. Panel C of Table 7 presents the results, which remain consistent with our baseline findings. 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.

7.4 Excluding AI patents

Our innovation measures used in the baseline analysis are constructed using utility patents granted by the USPTO, including AI patents that explicitly utilise AI-related technologies. According to the classification of Giczy et al. (2022), AI patents are those that contain at least one of the following technological components: machine learning, natural language processing, computer vision, speech, knowledge processing, AI hardware, evolutionary computation, and planning and control. If AI adoption has a profound impact on firm-level innovation beyond AI-specific technologies, our results should remain consistent even when AI patents are excluded from the construction of innovation measures.

To examine whether the effect of AI adoption is limited to AI related invention, we exclude AI patents when constructing innovation measures and re-estimate Equation 6. As shown in Panel D of Table 7, the results remain consistent with our baseline findings. This robustness check supports the interpretation that AI adoption enhances broader corporate innovation capabilities, rather than merely stimulating activity within AI-specific technological domains.

7.5 Excluding high-tech firms

The results in Section 6.5 indicate that, in general, there are no significant differences in the impact of AI on innovation between firms in high-tech and non-high-tech industries. These results are consistent with AI generating meaningful innovation gains outside high-tech industries, rather than its impact being concentrated exclusively in high-tech firms. To further assess whether the effects of AI on innovation are driven by technology-intensive industries, we exclude firms that operate in tech sectors in 2015 from our baseline sample and re-estimate Equation 6.

Table 7 Panel E shows that the positive effect of AI on innovation remains significant after excluding high-tech firms. Our results are consistent with Babina et al. (2024), which document that the benefits of AI adoption span a wide range of economic sectors, not just

those traditionally classified as high-tech.

8 Conclusions

This paper investigates the impact of AI adoption on corporate innovation in terms of quantity, quality, and strategy. Using a stacked DiD design, we find that AI adoption increases both the quantity and quality of innovation. AI-adopting firms produce more patents and their post-adoption patents receive more citations relative to non-adopting firms. AI adoption also reshapes firms' innovation strategies. The patents of AI-adopting firms are more original and general, include more claims, and reach more technologically distant classes. Adopters also steer towards more exploitative innovation. The results indicate that AI-adopting firms expand search into technologically distant domains while simultaneously build upon their existing technological capabilities. Our results remain robust after accounting for potential biases associated with endogenous AI adoption using a control-function IV based on pre-2010 exposure to AI-strong universities.

We also find that the impact of AI is more pronounced among firms with a focused business scope, consistent with the idea that AI technologies achieve better performance when fueled by high-quality and consistent data. Additionally, our results suggest that firms operating in highly competitive product markets leverage AI to generate innovation outputs that are not only more original but also more grounded in their existing technological strengths. This supports the view that firms under intense product market competition leverage AI to generate more differentiated innovation outputs and achieve incremental improvements with greater efficiency.

Our results are robust to entropy balancing, placebo tests, alternative control groups, and are not driven by large firms, financial constraints, or high-tech industries. In addition, we document that the impact of AI on innovation extends beyond AI-specific technologies, with significant effects across a broad range of economic sectors. Overall, our findings suggest that

AI enhances firms' ability to process and apply knowledge, thereby increasing the quantity, improving the quality, and refining the strategic orientation of corporate innovation.

REFERENCES

Accenture, "Reinventing R&D in the Age of AI," https://www.accenture.com/content/dam/accenture/final/accenture-com/document-2/Reinventing-RandD-In-The-Age-Of-AI-Report.pdf 2023.

Acemoglu, Daron, "The Simple Macroeconomics of AI," *Economic Policy*, 2025, *40* (121), 13–58.

Acharya, Viral and Zhaoxia Xu, "Financial Dependence and Innovation: The Case of Public versus Private Firms," *Journal of Financial Economics*, 2017, *124* (2), 223–243.

Aghion, Philippe, Ufuk Akcigit, Antonin Bergeaud, Richard Blundell, and David Hémous, "Innovation and Top Income Inequality," *The Review of Economic Studies*, 2019, *86* (1), 1–45.

Agrawal, Ajay, John McHale, and Alexander Oettl, "Finding Needles in Haystacks: Artificial Intelligence and Recombinant Growth," in "The economics of artificial intelligence: An agenda," University of Chicago Press, 2019, pp. 149–174.

_ , _ , and _ , "Artificial intelligence and scientific discovery: A model of prioritized search," *Research Policy*, 2024, *53* (5), 104989.

_ , Joshua Gans, and Avi Goldfarb, "Prediction, Judgment, and Complexity: A Theory of Decision-making and Artificial Intelligence," in "The Economics of Artificial Intelligence: An Agenda," University of Chicago Press, 2018, pp. 89–110.

Akcigit, Ufuk, Murat Alp Celik, and Jeremy Greenwood, "Buy, Keep, or Sell: Economic Growth and the Market for Ideas," *Econometrica*, 2016, *84* (3), 943–984.

_ , Salomé Baslandze, and Stefanie Stantcheva, "Taxation and the International Mobility of Inventors," *American Economic Review*, 2016, *106* (10), 2930–2981.

Alderucci, Dean and Douglas Sicker, "Applying Artificial Intelligence to The Patent System," *Technology & Innovation*, 2019, *20* (4), 415–425.

_ , Lee Branstetter, Eduard Hovy, Andrew Runge, and Nikolas Zolas, "Quantifying the Impact of AI on Productivity and Labor Demand: Evidence from U.S. Census Microdata," in "Allied Social Science Associations—ASSA 2020 Annual Meeting" 2020.

Almeida, Heitor and Murillo Campello, "Financial Constraints, Asset Tangibility, and Corporate Investment," *The Review of Financial Studies*, 2007, *20* (5), 1429–1460.

Arts, Sam, Bruno Cassiman, and Jianan Hou, "Technology Differentiation, Product Market Rivalry and MA Transactions," *Strategic Management Journal*, 2025, *46* (4), 837–862.

Autor, David H, Frank Levy, and Richard J Murnane, "The Skill Content of Recent Technological Change: An Empirical Exploration," *The Quarterly Journal of Economics*, 2003, *118* (4), 1279–1333.

Azoulay, Pierre, Joshua Krieger, and Abhishek Nagaraj, "Old Moats for New Models: Openness, Control, and Competition in Generative Artificial Intelligence," *Entrepreneurship and Innovation Policy and the Economy*, 2025, *4* (1), 7–46.

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson, "Artificial Intelligence, Firm Growth, and Product Innovation," *Journal of Financial Economics*, 2024, *151*, 103745.

_ , _ , Alex X He, and James Hodson, "Firm Investments in Artificial Intelligence Technologies and Changes in Workforce Composition," *National Bureau of Economic Research Working Paper*, 2023, *31325*.

Baker, Andrew C, David F Larcker, and Charles CY Wang, "How Much Should We Trust Staggered Difference-in-Differences Estimates?," *Journal of Financial Economics*, 2022, *144* (2), 370–395.

Begenau, Juliane, Maryam Farboodi, and Laura Veldkamp, "Big Data in Finance and the Growth of Large Firms," *Journal of Monetary Economics*, 2018, *97*, 71–87.

Benner, Mary J and Michael Tushman, "Process Management and Technological Innovation: A Longitudinal Study of the Photography and Paint Industries," *Administrative Science Quarterly*, 2002, *47* (4), 676–707.

Bennett, Nathaniel R, Joseph L Watson, Robert J Ragotte, Andrew J Borst, DéJenaé L See, Connor Weidle, Riti Biswas, Yutong Yu, Ellen L Shrock, Russell Ault et al., "Atomically accurate de novo design of antibodies with RFdiffusion," *Nature*, 2025, pp. 1–11.

Bresnahan, Timothy F. and Manuel Trajtenberg, "General Purpose Technologies 'Engines of Growth'?," *The Quarterly Journal of Economics*, 2017, *132* (2), 665–712.

Bresnahan, Timothy F, Erik Brynjolfsson, and Lorin M Hitt, "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-level Evidence," *The Quarterly Journal of Economics*, 2002, *117* (1), 339–376.

Brynjolfsson, Erik and Andrew McAfee, "What's Driving the Machine Learning Explosion?," *Harvard Business Review*, 2017, *18* (3), 118.

_ , Daniel Rock, and Chad Syverson, "Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics," in "The economics of artificiatificial intelligence: An agenda," University of Chicago Press, 2019, pp. 23–60.

_ , _ , and _ , "The Productivity J-curve: How Intangibles Complement General Purpose Technologies," *American Economic Journal: Macroeconomics*, 2021, *13* (1), 333–372.

_ , Danielle Li, and Lindsey Raymond, "Generative AI at Work," *The Quarterly Journal of Economics*, 2025, *140* (2), 889–942.

Cao, Sean, Wei Jiang, Baozhong Yang, and Alan L Zhang, "How to Talk When a Machine Is Listening: Corporate Disclosure in the Age of AI," *The Review of Financial Studies*, 2023, *36* (9), 3603–3642.

Chen, Mark A and Joanna Xiaoyu Wang, "Displacement or Augmentation? The Effects of AI on Workforce Dynamics and Firm Value," *SSRN Working Paper*, 2024.

Chu, Uyen, "AI in R&D: Top Use Cases You Need To Know," https://smartdev.com/ai-use-cases-in-research-and-development/ 2025. SmartDev.

Cockburn, Iain M, Rebecca Henderson, and Scott Stern, *The Impact of Artificial Intelligence on Innovation: An Exploratory Analysis*, Vol. 24449, National Bureau of Economic Research Cambridge, MA, USA, 2018.

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

Comment, Robert and Gregg A Jarrell, "Corporate Focus and Stock Returns," *Journal of Financial Economics*, 1995, *37* (1), 67–87.

Damioli, Giacomo, Vincent Van Roy, and Daniel Vertesy, "The Impact of Artificial Intelligence on Labor Productivity," *Eurasian Business Review*, 2021, *11*, 1–25.

Doshi, Anil R and Oliver P Hauser, "Generative AI Enhances Individual Creativity but Reduces the Collective Diversity of Novel Content," *Science Advances*, 2024, *10* (28), eadn5290.

Eisfeldt, Andrea L, Gregor Schubert, and Miao Ben Zhang, "Generative AI and firm values," *Journal of Finance*, 2025, *Forthcoming*.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock, "GPTs Are GPTs: Labor Market Impact Potential of LLMs," *Science*, 2024, *384* (6702), 1306–1308.

Felten, Edward, Manav Raj, and Robert Seamans, "Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses," *Strategic Management Journal*, 2021, *42* (12), 2195–2217.

Fitzgerald, Tristan and Xiaoding Liu, "Shared Culture and Technological Innovation: Evidence From Corporate R&D Teams," *SSRN Working Paper*, 2020.

Fleming, Lee, "Recombinant Uncertainty in Technological Search," *Management science*, 2001, *47* (1), 117–132.

Fujii, Hidemichi and Shunsuke Managi, "Trends and Priority Shifts in Artificial Intelligence Technology Invention: A Global Patent Analysis," *Economic Analysis and Policy*, 2018, *58*, 60–69.

Gao, Huasheng, Po-Hsuan Hsu, and Kai Li, "Innovation Strategy of Private Firms," *Journal of Financial and Quantitative Analysis*, 2018, *53* (1), 1–32.

Gao, Wenlian and Julia Chou, "Innovation Efficiency, Global Diversification, and Firm Value," *Journal of Corporate Finance*, 2015, *30*, 278–298.

Giczy, Alexander V, Nicholas A Pairolero, and Andrew A Toole, "Identifying Artificial Intelligence (AI) Invention: A Novel AI Patent Dataset," *The Journal of Technology Transfer*, 2022, *47* (2), 476–505.

Gormley, Todd A and David A Matsa, "Growing Out of Trouble? Corporate Responses to Liability Risk," *The Review of Financial Studies*, 2011, *24* (8), 2781–2821.

__ and __, "Playing It Safe? Managerial Preferences, Risk, and Agency Conflicts," *Journal of Financial Economics*, 2016, *122* (3), 431–455.

Griliches, Zvi, "Patent Statistics as Economic Indicators: A Survey," in "R&D and productivity: the econometric evidence," University of Chicago Press, 1998, pp. 287–343.

Haefner, Naomi, Joakim Wincent, Vinit Parida, and Oliver Gassmann, "Artificial Intelligence and Innovation Management: A Review, Framework, and Research Agenda," *Technological Forecasting and Social Change*, 2021, *162*, 120392.

Hainmueller, Jens, "Entropy Balancing for Causal Effects: A Multivariate Reweighting Method to Produce Balanced Samples in Observational Studies," *Political Analysis*, 2012, *20* (1), 25–46.

Hall, Bronwyn H, Adam B Jaffe, and Manuel Trajtenberg, "The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools," 2001.

Hall, Bronwyn H., Adam Jaffe, and Manuel Trajtenberg, "Market Value and Patent Citations," *The RAND Journal of Economics*, 2005, *36* (1), 16–38.

Hall, Bronwyn H, Zvi Griliches, and Jerry A Hausman, "Patents and R&D: Is There a Lag?," *National Bureau of Economic Research Working Paper*, 1984.

He, Qiyang and Buhui Qiu, "Environmental Enforcement Actions and Corporate Green Innovation," *Journal of Corporate Finance*, 2025, *91*, 102711.

Hoberg, Gerard and Gordon Phillips, "Text-Based Network Industries and Endogenous Product Differentiation," *Journal of Political Economy*, 2016, *124* (5), 1423–1465.

__ **and Vojislav Maksimovic**, "Redefining Financial Constraints: A Text-Based Analysis," *The Review of Financial Studies*, 2015, *28* (5), 1312–1352.

Jansen, Justin JP, Frans AJ Van Den Bosch, and Henk W Volberda, "Exploratory Innovation, Exploitative Innovation, and Performance: Effects of Organizational Antecedents and Environmental Moderators," *Management Science*, 2006, *52* (11), 1661–1674.

Jiang, Wei, Junyoung Park, Rachel Jiqiu Xiao, and Shen Zhang, "AI and the Extended Workday: Productivity, Contracting Efficiency, and Distribution of Rents," *National Bureau of Economic Research Working Paper*, 2025.

Jones, Benjamin F, "The Burden of Knowledge and the "Death of the Renaissance Man": Is Innovation Getting Harder?," *The Review of Economic Studies*, 2009, *76* (1), 283–317.

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

Lee, Peter, "Patent Law and the Two Cultures," *Yale Law Journal*, 2010, *120*, 2.

Leone, Maria Isabella and Toke Reichstein, "Licensing-in Fosters Rapid Invention! The Effect of the Grant-back Clause and Technological Unfamiliarity," *Strategic Management Journal*, 2012, *33* (8), 965–985.

Lerner, Josh and Amit Seru, "The Use and Misuse of Patent Data: Issues for Finance and Beyond," *The Review of Financial Studies*, 2022, *35* (6), 2667–2704.

Lerner, Joshua, "The Importance of Patent Scope: an Empirical Analysis," *The RAND Journal of Economics*, 1994, pp. 319–333.

Li, Dongmei, "Financial Constraints, R&D Investment, and Stock Returns," *The Review of Financial Studies*, 2011, *24* (9), 2974–3007.

Li, Kai and Jin Wang, "Inter-Firm Inventor Collaboration and Path-Breaking Innovation: Evidence From Inventor Teams Post-Merger," *Journal of Financial and Quantitative Analysis*, 2023, *58* (3), 1144–1171.

Lu, Yao, Gordon M Phillips, and Jia Yang, "The Impact of Cloud Computing and AI on Industry Dynamics and Concentration," *National Bureau of Economic Research Working Paper*, 2024.

Ludwig, Jens and Sendhil Mullainathan, "Machine Learning as a Tool for Hypothesis Generation," *The Quarterly Journal of Economics*, 2024, *139* (2), 751–827.

Ma, Song, Joy Tianjiao Tong, and Wei Wang, "Bankrupt Innovative Firms," *Management Science*, 2022, *68* (9), 6971–6992.

McKinsey & Company, "The Next Innovation Revolution Powered by AI," https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-next-innovation-revolution-powered-by-ai 2023.

McPhee, Devon, "AI Algorithm Identifies High-Performing Electrolytes for Batteries," https://www.mccormick.northwestern.edu/news/articles/2024/06/ai-algorithm-identifies-high-performing-electrolytes-for-batteries/ 2024. Northwestern University Engineering News.

Merchant, Amil, Simon Batzner, Samuel S Schoenholz, Muratahan Aykol, Gowoon Cheon, and Ekin Dogus Cubuk, "Scaling Deep Learning for Materials Discovery," *Nature*, 2023, *624* (7990), 80–85.

Microsoft, "Unlocking a New Era for Scientific Discovery with AI: How Microsoft's AI Screened Over 32 Million Candidates to Find a Better Battery," https://azure.microsoft.com/en-us/blog/quantum/2024/01/09/unlocking-a-new-era-for-scientific-discovery-with-ai-how-microsofts-ai-screened-over-32-million-candidates-to-find-a-better-battery/ 2024.

Mihet, Roxana and Thomas Philippon, "The Economics of Big Data and Artificial Intelligence," in "Disruptive Innovation in Business and Finance in the Digital World," Emerald Publishing Limited, 2019, pp. 29–43.

Moretti, Enrico, "The Effect of High-tech Clusters on the Productivity of Top Inventors," *American Economic Review*, 2021, *111* (10), 3328–3375.

Mühlroth, Christian and Michael Grottke, "Artificial Intelligence in Innovation: How to Spot Emerging Trends and Technologies," *IEEE Transactions on Engineering Management*, 2020, *69* (2), 493–510.

Ouellette, Lisa Larrimore, Victoria Fang, and Nicholas T. Ouellette, "How Will AI Affect Patent Disclosures?," *Nature Biotechnology*, 2025, *43*, 26–28.

Peters, Ryan H and Lucian A Taylor, "Intangible capital and the investment-q relation," *Journal of Financial Economics*, 2017, *123* (2), 251–272.

Rajan, Raghuram, Henri Servaes, and Luigi Zingales, "The Cost of Diversity: The Diversification Discount and Inefficient Investment," *The Journal of Finance*, 2000, *55* (1), 35–80.

Setchi, Rossitza, Irena Spasić, Jeffrey Morgan, Christopher Harrison, and Richard Corken, "Artificial Intelligence for Patent Prior Art Searching," *World Patent Information*, 2021, *64*, 102021.

Stoffman, Noah, Michael Woeppel, and M Deniz Yavuz, "Small Innovators: No Risk, No Return," *Journal of Accounting and Economics*, 2022, *74* (1), 101492.

Trajtenberg, Manuel, Rebecca Henderson, and Adam Jaffe, "University Versus Corporate Patents: A Window On The Basicness Of Invention," *Economics of Innovation and New Technology*, 1997, *5* (1), 19–50.

Tranchero, Matteo, "Finding Diamonds in the Rough: Data-Driven Opportunities and Pharmaceutical Innovation," *UC Berkeley Working Paper*, 2024.

Wamba-Taguimdje, Serge-Lopez, Samuel Fosso Wamba, Jean Robert Kala Kamdjoug, and Chris Emmanuel Tchatchouang Wanko, "Influence of Artificial Intelligence (AI) on Firm Performance: The Business Value of AI-Based Transformation Projects," *Business Process Management Journal*, 2020, *26* (7), 1893–1924.

Webb, Michael, "The Impact of Artificial Intelligence on the Labor Market," *SSRN Working Paper*, 2019.

__ , Nick Short, Nicholas Bloom, and Josh Lerner, "Some Facts of High-tech Patenting," 2018.

Weitzman, Martin L, "Recombinant Growth," *The Guarterly Journal of Economics*, 1998, *113* (2), 331–360.

Whang, Steven Euijong, Yuji Roh, Hwanjun Song, and Jae-Gil Lee, "Data Collection and Quality Challenges in Deep Learning: A Data-Centric AI Perspective," *The VLDB Journal*, 2023, *32* (4), 791–813.

Wills, Todd, "AI drug discovery: assessing the first AI-designed drug candidates to go into human clinical trials," CAS Insights 2022.

Wooldridge, Jeffrey M, "Control Function Methods in Applied Econometrics," *Journal of Human Resources*, 2015, *50* (2), 420–445.

Wuchty, Stefan, Benjamin F Jones, and Brian Uzzi, "The Increasing Dominance of Teams in Production of Knowledge," *Science*, 2007, *316* (5827), 1036–1039.

Yoo, Hyoung Sun, Ye Lim Jung, and Seung-Pyo Jun, "Prediction of SMEs' R&D Performances by Machine Learning for Project Selection," *Scientific Reports*, 2023, *13* (1), 7598.

Yordy, Lucas R., "The Library of Babel for Prior Art: Using Artificial Intelligence to Mass Produce Prior Art in Patent Law," *Vanderbilt Law Review*, 2021, *74*, 521–562.

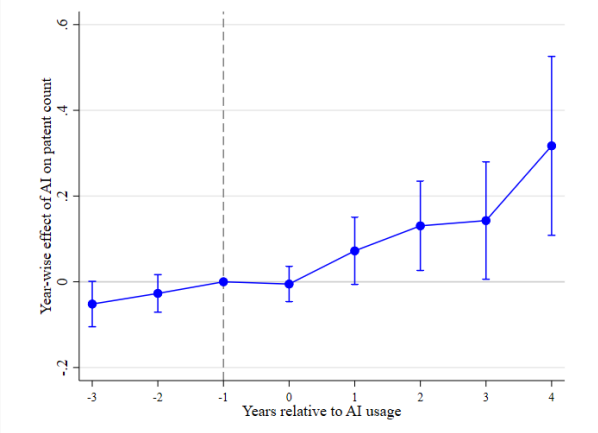
Zha, Daochen, Zaid Pervaiz Bhat, Kwei-Herng Lai, Fan Yang, Zhimeng Jiang, Shaochen Zhong, and Xia Hu, "Data-centric Artificial Intelligence: A Survey," *ACM Computing Surveys*, 2025, *57* (5), 1–42.

Zheng, Xiang, "How Can Innovation Screening Be Improved? A Machine Learning Analysis with Economic Consequences for Firm Performance," *Journal of Financial and Quantitative Analysis*, 2025, p. 1–57.

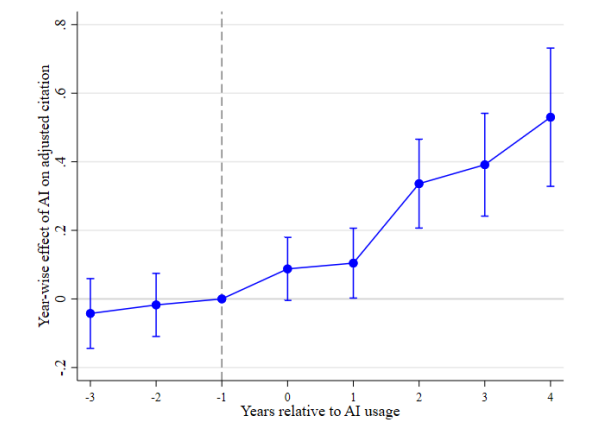
Figure 1
Dynamic Effects of AI Adoption

This figure reports the estimated coefficients from the lead-lag model, capturing the dynamic effects of AI adoption. The horizontal axis denotes years relative to AI adoption. The horizontal axis varies from year -3 to year 4, with year 0 representing the year of adoption within each cohort. We use year -1 as the reference point. Error bars indicate 95% confidence intervals. In Panel A, Figure (a) shows the dynamic treatment effects of AI adoption on innovation quantity measured by *patent count*. Figure (b) show the dynamic treatment effects of AI adoption on innovation quality measured by *adjusted citation*. Panel B, C, and D show the dynamic treatment effects of AI adoption on innovation strategy in terms of innovation novelty, innovation direction, and innovation reach, respectively. In Panel B, Figure (c) and (d) present the dynamic treatment effects of AI adoption on innovation novelty measured by *originality* and *distance*. In Panel C, Figure (e) and (f) present the dynamic treatment effects of AI adoption on innovation direction measured by *exploitative* and *explorative*. In Panel D, Figure (g) and (h) present the dynamic treatment effects of AI adoption on innovation reach measured by *generality* and *claim*. In all panels, we use stacked difference-in-differences (DiD) sample and apply the Pseudo-Poisson Maximum Likelihood estimation method. All specifications include Year-Cohort FE and Firm-Cohort FE. Standard errors are clustered at the firm level. Variable definitions and sources are provided in Appendix A.1.

Panel A: Innovation Quantity and Quality



(a) Patent count

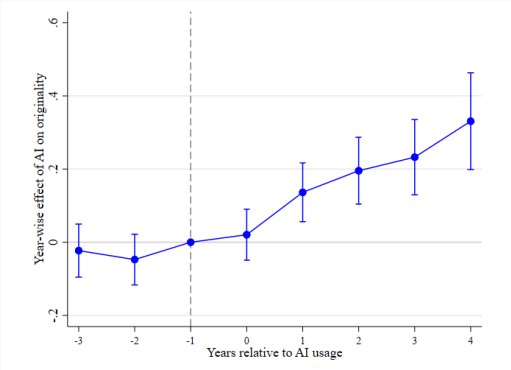


(b) Adjusted citation

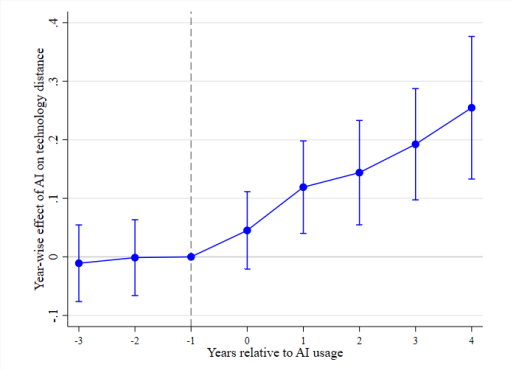
(continued on next page)

Figure 1 (continued)

Panel B: Innovation Novelty

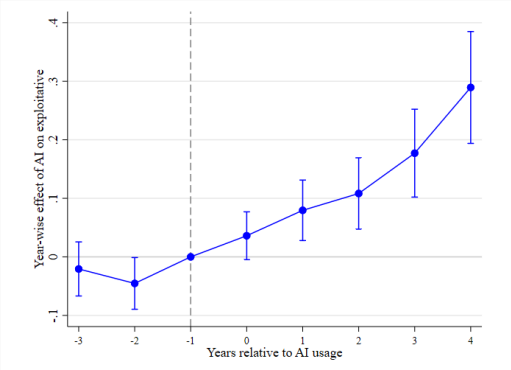


(c) Originality

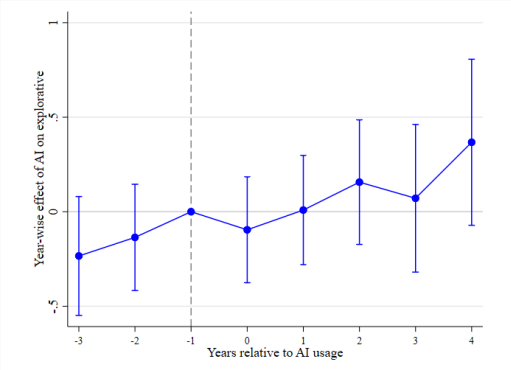


(d) Distance

Panel C: Innovation Direction

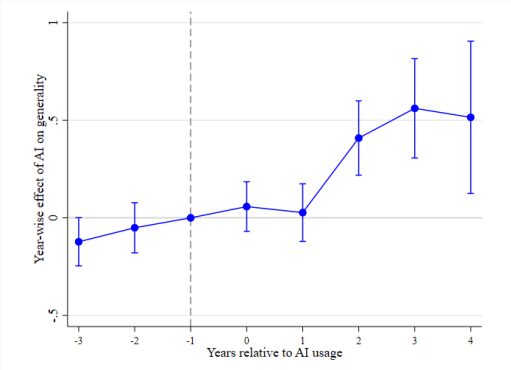


(e) Exploitative

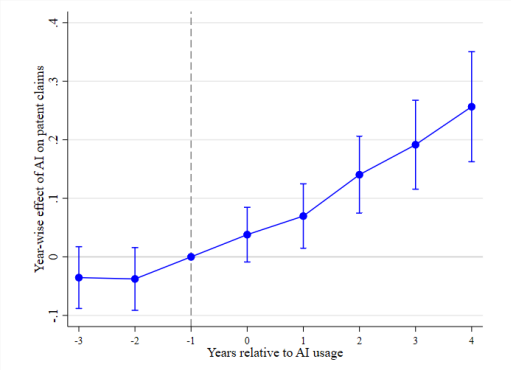


(f) Explorative

Panel D: Innovation Reach



(g) Generality



(h) Claim

Table 1
Summary Statistics

This table reports the summary statistics for all firms in the firm-year panel sample from 2011 to 2021. *All firms* refer to all unique firms in the sample. *AI-adopting firms* refer to firms that ever adopts AI technology from 2011 to 2021. *Never-adopting firm* refers to firms that never adopt AI technology from 2011 to 2021. Columns (1)-(3) report the summary statistics for all firms, columns (4)-(6) report the summary statistics for AI-adopting firms, and columns (7)-(10) report the summary statistics for non-adopting firms. Variable definitions and sources are provided in Appendix Table A.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
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	2.389	1.721	1.876
R&D	0.078	0.021	0.127	0.047	0.017	0.071	0.090	0.023	0.141

Table 2
AI and Innovation: Quantity and Quality

This table presents results from stacked DiD regressions estimating the impact of AI adoption on innovation outcomes. All models are estimated using Pseudo-Poisson Maximum Likelihood. Panel A presents results on patent quantity measured by *patent count* from *t+1* to *t+3* years following AI adoption in year *t*. Panel B presents results on patent quality measured by *adjusted citation* from *t+1* to *t+3* years following AI adoption in year *t*. *AI* is a dummy 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 A.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-squared	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-squared	0.261	0.255	0.248
Observations	53,623	42,190	32,551

Table 3
AI and Innovation: Strategy

This table presents results from stacked DiD regressions estimating the impact of AI adoption on innovation outcomes. In Panel A, we present results on innovation novelty measured by *originality* and *distance* from *t+1* to *t+3* years following AI adoption in year *t*. In Panel B, we present results on innovation direction measured by *exploitative* and *explorative* from *t+1* to *t+3* years following AI adoption in year *t*. In Panel C, we present results on innovation reach measured by *generality* and *claim* from *t+1* to *t+3* years following AI adoption in year *t*. All models are estimated using Pseudo-Poisson Maximum Likelihood. *AI* is a dummy 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 A.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-squared	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-squared	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-squared	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 the results of Control Function (CF) regressions on the effect of AI on corporate innovation using the stacked DiD sample. The first-stage result on *AI exposure* is reported in column (1). The second-stage results on corporate innovation measures in year *t+3* following AI adoption in year *t* are reported in columns (2)-(9). Column (2) presents results on innovation quantity, column (3) presents results on innovation quality, and columns (4)-(9) present results on innovation strategy. We estimate the probit model in the first stage. In the second stage, we use Pseudo-Poisson Maximum Likelihood method. *AI exposure* is 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* 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 A.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***								
	(10.37)								
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**
	(-2.61)	(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
Likelihood Ratio Chi-squared	1642.36								
Pseudo R-squared	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 *inventor pool*, *first-time inventor*, *AI-experienced inventor*, and *specialization*. Panel C reports results for the knowledge recombination channel, measured by *team size*, *unique inventor team*, *core skill breadth*, and *new cite ratio*. In all panels, we present Pseudo-Poisson Maximum Likelihood results from *t+1* to *t+3* years following AI adoption in year *t*. *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 A.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-squared	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												
	Inventor pool			First-time inventor			AI-experienced inventor			Specialization		
Variable	t+1	t+2	t+3	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)	(10)	(11)	(12)
AI	0.077*	0.094**	0.127***	0.164***	0.214***	0.199***	0.064*	0.070*	0.037	0.151***	0.147***	0.149***
	(1.65)	(2.16)	(3.43)	(2.93)	(3.52)	(2.80)	(1.69)	(1.75)	(0.87)	(6.61)	(5.56)	(5.14)
Year-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Pseudo R-squared	0.948	0.952	0.956	(2.93)	(3.52)	(2.80)	(1.69)	(1.75)	(0.87)	(6.61)	(5.56)	(5.14)
Observations	60,738	48,606	38,245	51,044	40,567	31,523	46,628	37,336	29,300	59,355	47,522	37,417
Panel C: Knowledge Recombination												
	Team size			Unique inventor team			Core skill breadth			New cite ratio		
Variable	t+1	t+2	t+3	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)	(10)	(11)	(12)
AI	0.142***	0.150***	0.154***	0.080*	0.106**	0.159***	0.136***	0.144***	0.142***	0.159***	0.142***	0.165***
	(5.69)	(5.43)	(5.08)	(1.67)	(2.35)	(4.17)	(5.75)	(5.44)	(5.00)	(4.67)	(3.78)	(4.01)
Year-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Pseudo R-squared	0.237	0.231	0.227	0.939	0.944	0.948	0.142	0.138	0.133	0.125	0.125	0.123
Observations	60,738	48,606	38,245	60,738	48,606	38,245	59,355	47,522	37,417	57,584	45,973	35,981

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 corporate innovation. Panel A reports the results on firms with different business focus by comparing the differential effects between focused and diversified firms. Panel B presents the results on firms facing different product market competition by comparing the differential effects between competitive and non-competitive firms. Panel C reports the results on firms of different sizes by comparing the differential effects between big and small firms. Panel D presents the results on firms facing different financial constraints by comparing the differential effects between financially-constrained and other firms. Panel E reports the results on firms with different technological capabilities by comparing the differential effects between high-tech and other firms. In all panels, we present Pseudo-Poisson Maximum Likelihood results in year *t+3* following AI adoption in year *t*. *AI* is an indicator variable that equals 1 if firm *i* in cohort *c* uses AI technology in year *t*, and 0 otherwise. In Panel A, all specifications include Year-Cohort FE, Firm-Cohort FE, Year-Cohort-Focused firm FE, and Firm-Cohort-Focused firm FE. In Panel B, all specifications include Year-Cohort FE, Firm-Cohort FE, Year-Cohort-High competition FE, and Firm-Cohort-High competition FE. In Panel C, all specifications include Year-Cohort FE, Firm-Cohort FE, Year-Cohort-Big firm FE, and Firm-Cohort-Big firm FE. In Panel D, all specifications include Year-Cohort FE, Firm-Cohort FE, Year-Cohort-Constrained firm FE, and Firm-Cohort-Constrained firm FE. In Panel E, all specifications include Year-Cohort FE, Firm-Cohort FE, Year-Cohort-High-Tech firm FE, and Firm-Cohort-High-Tech firm 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 A.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-squared	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

(continued on next page)

Table 6 (continued)

Panel B: Product Market Competition								
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-squared	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
Panel C: 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-squared	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 D: Financial Constraints								
Variable	Patent count	Adjusted citation	Originality	Distance	Exploitative	Explorative	Generality	Claim
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AI	0.157***	0.238***	0.183***	0.090***	0.132***	0.203	0.373***	0.133***
	(3.59)	(3.62)	(3.99)	(2.11)	(4.11)	(1.62)	(3.86)	(3.98)
AI x Constrained firm	0.064	0.314**	0.051	0.143*	0.084	0.059	0.270	0.080
	(0.70)	(2.54)	(0.63)	(1.84)	(1.35)	(0.24)	(1.35)	(1.27)
Year-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-Cohort-Constrained firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Cohort-Constrained firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Pseudo R-squared	0.938	0.251	0.0811	0.118	0.106	0.150	0.108	0.301
Observations	36,792	39,862	34,292	32,743	36,492	29,465	22,786	36,762
Panel E: 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-squared	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
AI and Innovation: Robustness Tests

This table presents the results of stacked DiD regressions on the effect of AI on corporate innovation in different robustness tests. Panel A reports the results of stacked DiD regressions 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 stacked DiD regressions from 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 the results of stacked DiD regressions using the never-treated firms (i.e., firms that never adopt AI technology from 2011 to 2021) as the control group. Panel D reports the stacked DiD estimates of the effect of AI on non-AI innovation. We exclude AI patents identified by Giczy et al. (2022) from the patent-level sample and re-construct corporate innovation measures without AI patents. Panel E presents the stacked DiD estimates of the effect of AI on corporate 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). All models are estimated using Pseudo-Poisson Maximum Likelihood. In all panels, we present the results on corporate innovation measures in year *t+3* following AI adoption in year *t*. Column (1) presents the results on innovation quantity, column (2) presents the results on innovation quality, and columns (4)-(9) present results on innovation strategy. *AI* is a dummy 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. Specifications in Panel E further include Year-Cohort-High-tech firm FE and Firm-Cohort-High-tech firm 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 A.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)	(2.19)	(2.71)	(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

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-squared	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

(continued on next page)

Table 7 (continued)

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-squared	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-squared	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
Panel E: Excluding High-tech Firms								
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-squared	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

Appendix A.

Table A.1
Variable Definition

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 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 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 A.1 (continued)

Variable	Description	Source
	B. AI variables	
AI	An indicator variable that equals 1 if firm i in cohort c uses AI technology in year t, and 0 otherwise.	SWZD
AI patent	Patent p filed by firm i in year t that is predicted by Giczy et al. (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).	Compustat
Leverage	The sum of long-term debt (dltt) and debt in current liabilities (dlc), divided by total assets (at).	Compustat
Cash	The sum of cash and short-term investments (che), divided by total assets (at).	Compustat
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).	Compustat
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 as 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 as $G_{i,t} = (1-\delta_{R\&D})G_{i,t-1} + R\&D_{i,t}$, 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 as $C_{i,t} = (1 - \delta_{SG\&A})C_{i,t-1} + \lambda SG\&A_{i,t}$, where $\delta_{SG\&A}$ is a depreciation rate of 20%, λ is a multiplier of 30%, the fraction of SG&A that is assumed to present an investment in intangible capital., and $SG\&A_{i,t}$ is selling, general and administrative expenses (xsga) minus R&D expenses (xrd).	Peters and Taylor (2017)

(continued on next page)

Table A.1 (continued)

Variable	Description	Source
	D. Financial variables (continued)	
R&D	R&D expenses (xrd) divided by total assets (at), or 0 if R&D expenses are negative or missing.	Compustat
Capital expenditure	Capital expenditures (capx) divided by total assets (at).	Compustat
Production cost	Cost of goods sold (cogs) divided by total sales (sale).	Compustat
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.	Compustat
Total operating cost	Total operating expenses (xopr) divided by total sales (sale).	Compustat
	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
Inventor pool	The total number of unique inventors who file at least one patent with 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
AI-experienced inventor	For firm i in year t, *AI-experienced 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 are AI-experienced. An inventor is classified as AI-experienced if she has filed at least one patent that is predicted by Giczy et al. (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 up to year *t-1*.	USPTO, Artificial Intelligence Patent Dataset

(continued on next page)

Table A.1 (continued)

Variable	Description	Source
	E. Channel variables (continued)	
Specialization	The average *Specialization* scores of all inventors who file at least one patent with firm *i* in year *t*. An inventor's Specialization score in year *t* is the Herfindahl 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 class.	USPTO
Team size	The average number of inventors per patent *p* filed by firm *i* in year *t*.	USPTO
Unique inventor team	The number of distinct inventor teams, made of unique combinations of inventors, that file at least one patent with firm *i* in year *t*.	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
New cite ratio	For firm *i* in year *t*, *new cite ratio* is the average of the patent-level new cite ratio across the firm's patents filed in year *t*. The new cite ratio for patent *p* filed by firm *i* in year *t* is defined as the share of backward citations made to new knowledge, that is, knowledge outside the inventor team's existing knowledge, relative to the total number of backward citations to U.S. patents, foreign patents, and other references (Fitzgerald and Liu, 2020). The inventor team's new knowledge consists of citations to U.S. patents, foreign patents, or other references that have neither been previously cited nor produced by any member of the inventor team prior to the filing date of patent *p*. The inventor team's existing knowledge comprises all U.S. patents, foreign patents, and other references that have been previously cited or produced by at least one team member before the filing date of patent *p*.	USPTO

(continued on next page)

Table A.1 (continued)

Variable	Description	Source
	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
High competition	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).	Hoberg and Phillips (2016)
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.	Compustat
Constrained firm	An indicator variable that equals 1 if the risk of firm *i* to delay investments due to issues with liquidity is equal to or larger than 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).	Compustat

Table A.2
Industry Distribution

This table presents the industry distribution of all firms in the firm-year panel sample from 2011 to 2021. *All firms* refer to all unique firms in the sample. *AI-adopting firms* refer to firms that ever adopts AI technology from 2011 to 2021. *Never-adopting firm* refers to firms that never adopt AI technology from 2011 to 2021. 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
Pharmaceutical Products	470	23	4.89	477	95.11
Computer Software	354	101	28.53	253	71.47
Electronic Equipment	218	42	19.27	176	80.73
Medical Equipment	169	20	11.83	149	88.17
Machinery	113	17	15.04	96	84.96
Business Services	95	28	29.47	67	70.53
Petroleum and Natural Gas	79	18	22.78	61	77.22
Retail	75	28	37.33	47	62.67
Chemicals	72	9	12.50	63	87.50
Computer Hardware	65	15	23.08	50	76.92
Measuring and Control Equipment	63	22	34.92	41	65.08
Communication	62	18	29.03	44	70.97
Wholesale	59	18	30.51	41	69.49
Automobiles and Trucks	58	12	20.69	46	79.31
Electrical Equipment	55	6	10.91	49	89.09
Utilities	51	18	35.29	33	64.71
Construction Materials	51	8	15.69	43	84.31
Transportation	44	20	45.45	24	54.55
Consumer Goods	40	11	27.50	29	72.50
Healthcare	32	9	28.12	23	71.88
Food Products	31	7	22.58	24	77.42
Construction	27	2	7.41	25	92.59
Steel Works Etc	26	4	15.38	22	84.62
Business Supplies	24	4	16.67	20	83.33
Recreation	23	5	21.74	18	78.26
Apparel	20	4	20.00	16	80.00
Aircraft	19	6	31.58	13	68.42
Entertainment	18	7	38.89	11	61.11
Almost Nothing	18	2	11.11	16	88.89
Personal Services	17	4	23.53	13	76.47
Restaurants, Hotels, Motels	14	4	28.57	10	71.43
Rubber and Plastic Products	14	2	14.29	12	85.71
Defense	10	2	20.00	8	80.00
Shipbuilding, Railroad Equipment	10	2	20.00	8	80.00
Non-Metallic and Industrial Metal Mining	10	0	0.00	10	100.00
Printing and Publishing	9	3	33.33	6	66.67
Shipping Containers	9	1	11.11	8	88.89
Tobacco Products	7	2	28.57	5	71.43
Agriculture	7	0	0.00	7	100.00
Beer and Liquor	6	1	16.67	5	83.33
Candy and Soda	5	1	20.00	4	80.00
Fabricated Products	4	0	0.00	4	100.00
Textiles	4	0	0.00	4	100.00
Coal	3	0	0.00	3	100.00
Precious Metals	3	0	0.00	3	100.00