S.T. Yau High School Science Award

Research Report

The Team

Name of team member: Wenzhao Liu

School: Phillips Exeter Academy

City, Country: Exeter, NH, USA

Name of supervising teacher: Shaocheng Ma

Job Title: Math Instructor

School/Institution: Phillips Exeter Academy

City, Country: Exeter, NH, USA

Title of Research Report

Cultivating the Soil or Breaking New Ground? Artificial Intelligence and Firms' Exploitative and Exploratory Innovation

Date

[September 10th]

Cultivating the Soil or Breaking New Ground? Artificial Intelligence and Firms' Exploitative and Exploratory Innovation

Wenzhao Liu

Abstract

How firms allocate resources between exploration-based innovation and exploratory innovation is the core proposition of organizational learning theory. The rapid penetration of artificial intelligence casts this classic problem in a new light: does AI encourage enterprises to work on their existing track, or does it lead them to unfamiliar fields? Drawing on a sample comprising publicly traded Chinese manufacturing firms over the period spanning 2010–2024, this paper develops a firm-specific index capturing AI application using textual information extracted from annual disclosures and measures innovation along two dimensions, patent type and knowledge distance. It first assesses how AI shapes overall innovative output and then tests whether this effect differs between innovation rooted in familiar knowledge and activity that explores new domains. The findings show that AI adoption significantly improved the overall volume of innovation generated by enterprises, but that its effect remains unevenly distributed. AI significantly promotes invention patents (breakthrough innovation) and exploitative innovation, whereas its effects on utility-model patents (incremental innovation) and exploratory innovation are insignificant. Artificial intelligence, therefore, drives enterprises to achieve high-quality breakthroughs along their existing technological track, presenting an innovative model of focus, deep cultivation, and quality improvement. Mechanism analysis indicates that artificial intelligence operates through adjustments in human capital composition, facilitating knowledge spillover, expanding knowledge diversity, improving the efficiency of capital allocation, and promoting a shift toward a lighter asset structure. Further comparisons reveal greater effects among large firms, firms in technology-intensive sectors, and firms facing intense product-market competition. The results offer empirical evidence on how artificial intelligence reshapes the innovation structure of enterprises and offer guidance for the allocation of innovation resources during firms' AI-enabled transformation.

Keywords: artificial intelligence; dual innovation; exploitative innovation; exploratory innovation; breakthrough innovation; incremental innovation

Acknowledgement

I would like to thank my instructor, Dr. Shaocheng MA (Phillips Exeter Academy, Math instructor), who provided me with the background knowledge and learning path of the subject and led me into the field of econometrics. In all aspects of scientific research, he answered my questions, taught me to think like a researcher, and laid the foundation for the success of my thesis.

1. Subject Source

The author's interest in both AI and economics goes back a long way. Over the past two years, AI has grown very fast and has drawn a great deal of attention, and more and more manufacturing firms are now using it. Economic and political news also began to attract closer attention. When the project started, the author picked up some basic econometrics and learned to run code in Stata. From the reading done for the project, and from everyday observation, AI seems to help firms do better with the products and technologies they already have, and to be less good at pushing them into brand-new technology areas. That observation led to the central question of this project: does AI mainly help firms "cultivate the soil they already own", or can it also help them "break new ground"?

Starting from that observation, the instructor suggested that the author look into how firms are affected by the use of AI. Introductory material was read, a few possible models were discussed with him, and a topic was eventually settled on: how the use of AI shapes the innovation structure of a firm. To study this at the firm level, the annual reports of listed companies were collected and the AI-related words in them were counted.

2. Relationship and Division of Labor Between the Instructor and the Student

Dr. Shaocheng Ma is the author's supervising teacher. His guidance was given free of charge, and there is no business relationship of any kind between them.

(1) Topic Selection and Research Design. The relevant literature was searched for and read independently, and the research questions were settled on with the instructor's advice. He also helped organize the logic of the theoretical framework and sharpen the ideas behind the hypotheses. The central parts of the research design, such as choosing variables and setting up the model, were completed by the author, with the instructor offering guidance along the way.

(2) Data Acquisition and Processing. The following tasks were carried out independently: collecting the annual reports of listed companies from http://www.cninfo.com.cn; obtaining patent data from the State Intellectual Property Office; downloading firm financial data from the CSMAR database; writing text-analysis code in Python to build an AI dictionary and count AI-related words in each annual report; and using Stata for data cleaning, variable construction and regression analysis. The instructor offered suggestions on how to standardize the data work and how to debug the code.

(3) Empirical Analysis. Every regression was run by the author: the baseline regression, the regression by innovation type, the instrumental-variable regression, the robustness checks, the mechanism analysis and the heterogeneity analysis. The instructor helped interpret the economic meaning of the results and suggested ways to improve how the tables are presented.

(4) Thesis Writing. The first draft was written independently. The instructor commented on the structure, on the logic of the argument and on the academic style, and the final version was produced after several rounds of revision.

3. Difficulties Encountered and Solutions

Difficulty 1: How to measure AI. At first the plan was simply to count how many

AI keywords appeared in an annual report. But annual reports differ a lot in length, and a raw count is strongly driven by that length. On the instructor's advice the logarithm of the AI word count was used instead, and a second measure was also built from the Management Discussion and Analysis (MD&A) section so that the two could be checked against each other.

Difficulty 2: Choosing an instrumental variable. Several candidates were tested before settling, with the instructor's help, on a Bartik-style instrument. It combines the industry-level AI intensity in 2009 (the initial share) with the national growth rate of the tertiary sector (the shock). The diagnostic tests showed that the instrument is not weak and that the model is properly identified.

Difficulty 3: Holding the argument together. In the first draft, the literature review read like a list of studies. The instructor suggested grouping the studies by theoretical theme instead, so that the review would set up the hypotheses rather than simply catalogue earlier work.

4. Whether Other Assistance is Paid

The data used here all come from public sources. Annual reports of listed companies were taken from http://www.cninfo.com.cn, patent records from the State Intellectual Property Office, firm financial data from the CSMAR database, and city-level figures from the statistical yearbook of Chinese cities. These are openly available research data, and they were used only to build variables and to run the empirical analysis reported here. Every source is named in the paper.

No laboratory equipment was needed, no unpublished results were used, and no calculations were performed by anyone else. All of the data cleaning, variable construction, regression work and robustness testing were carried out independently.

The instructor, Dr. Shaocheng Ma, gave academic guidance on the choice of topic, on the theoretical framework, on the measurement approach and on the revision of the paper. He was not paid for this, and there is no commercial relationship between the two parties.

No company, training provider or other commercial organization took part in this research or in the writing of this report. Apart from the public data listed above, the project received no outside funding and no commercial support.

5. Use of AI Tools

(1) AI Tools and Versions Used

With the instructor's permission and under his supervision, two AI tools were used as assistants. The first was ChatGPT (GPT-4o) and the second was Claude (3.5 Sonnet). They were used between June and August 2026, roughly once or twice a week.

(2) Specific Use Links and Purposes

1. Finding literature. While the literature review was being written, the AI tools were used to suggest extra keywords and to point towards possible papers. For instance, when work on "artificial intelligence and enterprise innovation" was being searched for, the tools listed the titles and authors of several papers that might be relevant. Every one of these was then checked in Google Scholar and in other academic databases, confirming that the papers exist and that anything quoted matches what the authors actually said. The AI tools were never used to invent references.

2. Debugging code. When Stata or Python threw an error while the data were being processed, the AI tools were used to help work out what had gone wrong. If, for example, merging two datasets produced a message about mismatched variable types, the tool would suggest a fix. The code was then edited and the fix was checked. The

final version of every piece of code was run and verified by the author.

3. Language polishing and grammar checks. After the abstract and the introduction had been translated into English, the AI tools were used to check the grammar and improve the wording, mainly looking at articles, tense consistency and how smoothly the sentences read. The edited text was then gone through sentence by sentence to make sure the meaning was still accurate and the style remained academic.

(3) Usage time and frequency

The tools were used between June and August 2026, on average once or twice a week, mainly for finding literature and for debugging code.

AI tools were not used to choose the topic or to write the paper. The essential parts of the work—choosing the topic, designing the research, building the core argument and writing it up—were all carried out independently. Anything the AI tools helped produce was reviewed, edited and checked before it was used.

Commitments on Academic Honesty and Integrity

We hereby declare that we

1. are fully committed to the principle of honesty, integrity and fair play throughout the competition.
2. actually perform the research work ourselves and thus truly understand the content of the work.
3. observe the common standard of academic integrity adopted by most journals and degree theses.
4. have declared all the assistance and contribution we have received from any personnel, agency, institution, etc. for the research work.
5. undertake to avoid getting in touch with assessment panel members in a way that may lead to direct or indirect conflict of interest.
6. undertake to avoid any interaction with assessment panel members that would undermine the neutrality of the panel member and fairness of the assessment process.
7. observe the safety regulations of the laboratory(ies) where we conduct the experiment(s), if applicable.
8. observe all rules and regulations of the competition.
9. agree that the decision of YHSA is final in all matters related to the competition.

We understand and agree that failure to honor the above commitments may lead to disqualification from the competition and/or removal of reward, if applicable; that any unethical deeds, if found, will be disclosed to the school principal of team member(s) and relevant parties if deemed necessary; and that the decision of YHSA is final and no appeal will be accepted.

(Signatures of full team below)



Name of team member: Wenzhao Liu



Name of supervising teacher: Shaocheng Ma

Table of Contents

1. Introduction 1
2. Literature Review and Theoretical Assumptions 2
3. Research Design 5
    3.1 Data Sources and Sample Construction 5
    3.2 Definition of Variables 5
    3.3 Specification of the Econometric Framework 7
4. Empirical Findings 7
    4.1 Baseline Regression Results 7
    4.2 Results by Innovation Category 8
    4.3 Addressing Endogeneity 9
    4.4 Robustness Tests 11
    4.5 Heterogeneity Analysis 12
5. Mechanism Analysis 13
    5.1 R&D Investment and Human Capital Upgrading 14
    5.2 Knowledge Spillover and Knowledge Diversity 14
    5.3 Resource Allocation Efficiency 14
6. Conclusion and Policy Implications 15
    6.1 Main Findings 15
    6.2 Policy Implications 16
References 18

1. Introduction

Does AI help a firm find things it does not yet know, or does it mostly help the firm get more out of what it already knows? The question matters because innovation is not one single activity. Firms can improve technologies they already use—exploitative innovation—or they can open up new technological paths—exploratory innovation (March, 1991). The first kind is about refinement, efficiency, and reliability. The second kind involves experiments, searching far from home, and returns that are hard to predict. Since both draw on the same limited attention and resources, a firm has to decide not only how much to innovate, but also which kind of innovation to go after.

AI could shift that decision. Because it can pick out patterns in very large datasets and can improve prediction and optimization, AI lowers the cost of searching inside a technology area the firm already knows. That is most useful when there is plenty of past data, when performance can be measured, and when the goal of the search is clear—exactly the conditions under which exploitative innovation thrives. Exploratory innovation is a harder case. Relevant data may be thin, the search space may be badly defined, and it may be impossible to judge in advance how valuable a discovery will be. So AI might do more to cultivate the "soil" a firm already owns than to help it break new ground. But that is not the only possibility. AI can also help a firm work with knowledge from outside its usual domains, spot connections nobody had noticed, and combine existing ideas in new ways. Which way the balance tips is therefore a central question to the study.

Evidence that AI raises innovation on average does not settle the question. Studies based on German, American, and Chinese data tend to find that firms that adopt or invest in AI innovate more (Rammer, Fernández, and Czarnitzki, 2022; Babina et al., 2024; Liu et al., 2020; Han and Mao, 2024; Jiang, Xuan, and Zhang, 2024). Other work shows that the effect varies with firm size, technology, and organizational capability (Grashof and Kopka, 2022; Brem, Giones, and Werle, 2023; Gao, Liu, and Yang, 2025). Together, these studies show that AI matters for innovation, but they say nothing about how it changes firms' strategies. A rise in total patenting could mean more work on familiar technologies, a move into unfamiliar ones, or both—and those two stories have very different consequences on firm growth, technological change, and how innovative capacity is distributed over the long run.

This paper fills that gap. It looks at AI use and patenting among Chinese manufacturing firms listed on the stock market between 2010 and 2024. A firm-level measure of how intensively a firm uses AI is built from what it discloses in its annual reports, and that measure is matched to detailed patent records. The patent data makes it possible to separate invention patents from utility-model patents, which represent, respectively, breakthrough innovation and incremental innovation. Meanwhile, through inspecting whether a patent falls in an International Patent Classification subclass the firm has already worked in or in one that lies outside its earlier experience (knowledge distance), this paper investigates AI's effect on both exploitative and exploratory innovation . This turns patenting activity into an empirical measure of exploitative and exploratory innovation. It then becomes possible to ask whether AI affects the two modes differently—and hence whether AI changes the mix of a firm's innovation portfolio.

The analysis goes further than simply recording an average relationship. It asks whether R&D spending, knowledge spillovers, and better resource allocation can

explain why AI affects innovation. These three channels capture three distinct ways the effect could work: by changing how many resources go into producing knowledge, by widening the knowledge a firm can draw on, and by improving how resources are spread across competing projects. The paper also looks at how the effects differ by firm size, ownership, and industry, which helps identify the organizational and technological settings where AI matters most.

This paper contributes to three strands of research. First, the analysis extends the study on AI-and-innovation away from merely the quantity of innovation outputs and towards the direction of it. The difference matters economically: a technology that makes existing research paths more productive need not make a firm any more willing or able to start new ones. By looking at exploitative and exploratory innovation separately, the paper shows whether AI widens a firm's technological frontier or mainly intensifies what happens inside the frontier it already has.

Second, the paper brings the technical nature of AI into the theory of organizational ambidexterity. That theory stresses the tension between exploitation, which depends on control and predictability, and exploration, which needs flexibility and a tolerance for uncertainty (Benner and Tushman, 2003). AI can shift this trade-off because it changes the cost of processing information, running experiments, and searching for knowledge. The framework used here therefore treats AI not as just one more input, but as a technology that can change the returns to different ways of learning within an organization.

Third, the paper supplies firm-level evidence on the mechanisms and the boundary conditions behind this reallocation. By looking at internal R&D resources, external knowledge spillovers, and resource-allocation efficiency side by side, it links a firm's adoption of AI to the process by which innovative opportunities get picked and developed. That helps explain not only whether AI promotes innovation, but also whether it pushes firms to deepen the capabilities they already have or to break new ground.

2. Literature Review and Theoretical Assumptions

AI is changing how firms innovate. Paschen, Pitt and Kietzmann (2020) classify the components of AI and contend that AI should be seen not merely as a device for raising efficiency but as a foundational technology capable of reshaping how firms generate and develop new ideas. Haefner et al. (2021) offer a broad framework showing the different roles AI plays as innovations are identified, developed and diffused. Mühlroth and Grottke (2022) note that AI helps firms spot emerging trends when the environment is uncertain. Verganti, Vendraminelli, and Iansiti (2020) stress that AI has altered how knowledge is produced during innovation and has opened new design options through collaboration between humans and machines. Hutchinson (2021) argues that AI's capability of innovating on its own may redraw the boundaries of innovation management, letting organizations take some innovation decisions without human involvement.

Numerous empirical analyses confirm that AI fosters innovation. Liu et al. (2020) use China's manufacturing panel data and document significant gains in technological innovation, whose magnitude varies from industry to industry. Han and Mao (2024) confirm this conclusion by using data from Chinese industrial enterprises. The research of Jiang, Xuan, and Zhang (2024) shows that AI adoption strengthens firms' innovation capacity, with larger gains among highly digitalized enterprises. Fu, Ni, and Fang (2025) discover that artificial intelligence has promoted the innovation of

digital enterprises through knowledge combination and process optimization. Ma, Shang, and Liang (2025) confirm complementarities between digital transformation and AI in shaping innovation outcomes. Gao, Liu, and Yang (2025) find that artificial intelligence enhances the perception, capture, and transformation capabilities at the firm level. Zhang and Li (2026) argue that AI can likewise stimulate innovative activity in small and medium-sized enterprises, although the outcome hinges on firms' absorptive capacity for new technologies. Ali and Frimpong (2020) discuss the scope for deploying AI in mining, noting that its innovative potential in resource-intensive industries has yet to be explored.

From a theoretical point of view, artificial intelligence influences innovation through multiple channels. Through a knowledge-management lens, artificial intelligence enhances firms' capacity to obtain, combine, and use knowledge (Li et al., 2025), and promotes innovation by broadening and deepening firms' knowledge search. In terms of resource allocation, artificial intelligence reconfigures production factors from inefficient processes to efficient innovation activities (Hassannataj Joloudari et al., 2024; Nagaranjan, 2022; Philip, 2024). Viewed through the lens of innovation management, artificial intelligence reduces the costs of experimentation and information processing, which in turn raises the efficiency of research and development (Haefner et al., 2021; Mühlroth and Grottke, 2022). The empirical record is broadly consistent with these arguments. Taken together, Liu et al. (2020), Han and Mao (2024), Jiang et al. (2024), Fu et al. (2025), Ma et al. (2025), Gao et al. (2025), and Zhang and Li (2026) provide consistent evidence of stronger innovation performance among firms with greater AI adoption. On this basis, this paper proposes:

Hypothesis H1: Firm-level AI adoption significantly increases aggregate innovation output.

However, AI need not affect all forms of innovation equally. Grashof and Kopka (2022) argue that firm size moderates AI's contribution to breakthrough innovation, with large enterprises benefiting more, which points to heterogeneous innovation responses to AI. Under the framework of dual innovation, Zhang et al. (2025) find that in China's artificial intelligence industry, exploratory innovation and exploitative innovation show completely different network characteristics and performance results. Zhao (2021) identifies knowledge-management ability and technological uncertainty as the main drivers of dual innovation. Audretsch, Belitski, and Fiedler (2025) revisit innovation and entrepreneurship through the lens of knowledge spillovers, arguing that spillovers form the link between AI and firm innovation. D'Amico et al. (2025) provide evidence that AI encourages innovation to appear and spread by making knowledge spillovers more efficient.

Exploitative innovation means optimizing and improving paths that already exist, and this fits well with what AI is good at—recognizing patterns and optimizing. AI performs best where data are plentiful and the problem is clearly defined, which lets it cut the cost of trial and error and make existing processes more efficient. Brem, Giones, and Werle (2023) show that AI has made the innovation-management process more systematic and therefore improved efficiency along established paths. Exploratory innovation is the opposite: it means breaking out of the existing framework and going into unknown territory, which calls for divergent thinking, a high tolerance for uncertainty, and the ability to decide when there are no data to go on. Choi et al. (2025) find that while AI helps create knowledge, it can also narrow the diversity of knowledge, which points to limits on tasks that need varied thinking. Zhang et al. (2025) likewise report a systematic difference between what drives

exploratory and exploitative innovation. Turning to patent types, invention patents face tougher examination and a higher inventive bar, and AI's strengths in integrating knowledge and recognizing patterns line up closely with what breakthrough innovation requires.

Paschen et al. (2020) find that AI affects innovation unevenly, noting that this heterogeneity is mirrored in the complexity of the innovations produced. Grashof and Kopka (2022) find that firm characteristics shape how AI affects breakthrough innovation. The following hypotheses follow:

Hypothesis H2a: AI has a more pronounced effect on exploitative innovation than on exploratory innovation.

Hypothesis H2b: AI affects breakthrough innovation more strongly than incremental innovation.

By what channels does AI's effect on innovation actually travel? AI works mainly by changing how a firm allocates its internal resources and how knowledge flows inside it. This paper, therefore, sets out a path that runs through three stages: input, process, and allocation.

On the input side, building, running, and maintaining AI systems calls for skilled people—people who know algorithms, data science, and software engineering (Xiao, Xiao, and He, 2024). Firms that adopt AI have to hire more R&D staff to keep up with digital transformation, which lifts their human capital towards higher skill levels (Verganti, Vendraminelli, and Iansiti, 2020; Hutchinson, 2021). Following the resource-based view, human capital with the right skills is the complementary asset that allows AI to have any effect. Therefore:

Hypothesis H3a: AI adoption fosters firm innovation by increasing the proportion of R&D personnel.

At the process stage, AI acts as an information-processing tool and changes how knowledge spreads and gets combined between firms. Because algorithms identify and match information, knowledge no longer moves only from one point to another; it circulates through a network in parallel, which strengthens knowledge externalities (D'Amico et al., 2025; Chin et al., 2024). Natural language processing and knowledge-graph techniques make it possible to capture the meaning of heterogeneous information and link it together in a structured way, which helps firms locate knowledge held outside their own boundaries (Audretsch, Belitski, and Fiedler, 2025). AI also helps firms combine knowledge more richly within fields they already work in, which widens their knowledge base. Accordingly:

Hypothesis H3b: AI use strengthens firm innovation by encouraging knowledge spillovers and knowledge diversity.

At the allocation stage, AI's ability to forecast lets a firm judge the return on an investment more accurately, which reduces the mismatch between capital and projects (Lu, Qin, and Yang, 2026). AI also links departments together, encourages collaboration and information sharing across them, and cuts duplicated work and wasted resources (Sjödin, Parida, and Kohtamäki, 2023). Because AI runs on software and algorithms, deploying it depends more on computing power, data, and talent than on physical fixed assets. That can loosen a firm's dependence on plant and equipment, freeing resources for intangible assets such as software, algorithms, and data and producing a leaner asset structure (Zhou, 2024). Therefore:

Hypothesis H3c: AI adoption enhances firm innovation by raising the efficiency with which resources are allocated.

AI's effect may also differ from one firm and one industry to the next. Large firms have more R&D money, better data infrastructure and more technical staff, all

of which sharpen their ability to recognize, absorb, and apply AI (Fonteneau et al., 2025). Grashof and Kopka (2022) report that the contribution of AI to breakthrough innovation is more pronounced in large firms. Zhang and Li (2026) show that how much small and medium-sized firms gain from AI depends on their capacity to absorb the technology. Therefore:

Hypothesis H4a: Enterprise scale positively regulates the innovation effect of AI;

Technology-intensive industries have a stronger technological base, greater research and development capabilities, and more highly-skilled workforce, which enables them to assimilate and apply artificial intelligence more effectively (Gao, Liu, and Yang, 2025). Lu, Du, and Qin (2022) evaluate national innovative capacity of 65 countries and conclude that countries with high technology intensity have significant advantages in dual innovation capabilities. Tang et al. (2020) find that the data reveal significant variation in the pace of AI-related innovation between different technical fields. Stinson and Vlaad (2024) point out that applications involving AI in high-technology domains are more diverse and creative. Therefore:

Assumption H4b: Industrial technology intensity positively regulates the innovation effect of artificial intelligence; that is, AI generates larger innovation gains in sectors with greater technological intensity.

The intensity of market competition shapes technology adoption and innovation incentives. Industrial organization theory holds that competitive pressure prompts enterprises to pursue technical progress to secure competitive advantage (Aghion et al., 2005). Enterprises operating in more competitive industries encounter stronger pressure to survive while seeking to compete for market share, which gives them a stronger motivation to apply artificial intelligence to product research and development and innovation. In less competitive industries, enterprises lack sufficient incentives for innovation, and artificial intelligence is more used for cost control than innovation. Therefore:

Hypothesis H4c: Competition within an industry amplifies the innovation effects of AI; in other words, artificial intelligence delivers a larger innovation payoff where competitive pressure is intense.

3. Research Design

3.1 Sample and Data

The dataset combines annual firm-level observations for manufacturing companies traded on China's A-share market during 2010–2024. The measure of AI adoption is derived from annual-report texts downloaded from the CNINFO platform. Patent information is matched from the China Economic and Financial Database (CCER), whereas firm characteristics and financial accounts are obtained from CSMAR. City characteristics are supplemented using the China Urban Statistical Yearbook. The raw sample is screened by removing financial firms, companies designated as ST, *ST, or PT, and observations lacking information on core variables. For each continuous variable, values below the 1st percentile or above the 99th percentile are replaced with the corresponding boundary values. After these procedures, the estimation sample contains 27,489 firm-year observations and forms an unbalanced panel.

3.2 Definition of Variables

1. Dependent Variables: Firm-Level Innovation

Innovation at the firm level is captured along two dimensions.

The first dimension concerns patent type. Because invention patents face more stringent examination and a higher inventiveness threshold, they serve as the proxy for breakthrough innovation. Utility-model patents, which carry lower inventiveness requirements, mainly capture incremental improvement. The variables are defined accordingly. Accordingly, the three outcomes are defined as the followings: ln_breakthrough = ln(1 + invention-patent applications), ln_incremental = ln(1 + utility-model patent applications), and ln_total = ln(1 + the combined number of invention and utility-model patent applications).

A second dimension is knowledge distance. This paper uses IPC four-digit subclass codes to define each firm's knowledge boundaries. A technology subclass is treated as familiar if the company has prior patent applications in that subclass. Exploitative innovation (ln_exploit) equals the natural logarithm of patent applications in familiar IPC subclasses plus one. Exploratory innovation (ln_explore) is the natural logarithm of patent applications in unfamiliar IPC subclasses plus one, calculated as total invention and utility-model applications minus those in familiar subclasses. Together, these two measures capture whether a firm is deepening its position in known technological areas or moving into new ones.

2. Main Explanatory Variable: Firm-Level AI Adoption (AI_apply)

The intensity of AI adoption is quantified at the firm level using textual information from corporate annual reports. First of all, the annual-report text for each listed firm is retrieved from Juchao Information Network before being organized into a firm-year panel data. Secondly, a dictionary of AI-related terms is constructed in three types: (1) determine the seed word set based on the classification of artificial intelligence technology and related industry research reports; (2) use Word2vec technology for semantic expansion; (3) add the expanded terminology dictionary to Python's Jieba sub-thesaurus; (4) sub-word the full text of each annual report Reason, and compute the frequency of accurate words and extended words after removing the disabled words. Finally, applying the transformation ln(1 + AI-related word count) yields a firm-specific annual artificial intelligence technology application index (AI_apply). Higher values of the index indicate more intensive applications of artificial intelligence technology.

3. Control Variables

Following relevant existing research, several covariates are selected along three dimensions: basic firm characteristics, financial condition, and corporate governance. Firm size (size) is defined as the natural logarithm of total assets and captures differences in the resources available for innovation. Firm age (lnage) is the logarithm of the number of years since establishment, accounting for experience accumulation and organizational inertia. Cash flow (cashflow), calculated as operating cash flow divided by total assets, proxies for the firm’s internal financing capacity. Leverage (lev) is measured as total liabilities relative to total assets and captures the extent to which debt obligations constrain innovation. Asset structure (asset), measured as the share of fixed assets, captures how asset allocation characteristics affect innovation flexibility. Total asset growth (growth) represents the growth of the enterprise's willingness to innovate, whereas return on assets (roa) controls for the role of profitability in securing internal innovation funding.

Table 1 summarizes the distributions of the variables used in the analysis. Among the dependent variables, overall innovation (ln_total) averages 2.675, and the average values of breakthrough innovation (ln_breakthrough) and incremental

innovation (ln_incremental) are 2.004 and 1.784 respectively. Exploitative innovation (ln_exploit) has a mean of 2.500, which is significantly higher than the mean of 0.988 for exploratory innovation (ln_explore), indicating that the sample firms concentrate their innovative activity largely within familiar technological fields. The key explanatory measure, AI_apply, averages 1.056 (SD = 1.330), revealing substantial dispersion in the intensity of AI use across firms. All control variables lie within reasonable ranges.

Table 1 Summary Statistics

VarName	Obs	Mean	SD	Min	Median	Max
ln_total	27489	2.675	1.227	0.693	2.565	6.394
ln_breakthrough	27489	2.004	1.246	0.000	1.792	5.861
ln_incremental	27489	1.784	1.386	0.000	1.792	5.576
ln_explore	27489	0.988	0.931	0.000	0.693	3.714
ln_exploit	27489	2.500	1.259	0.693	2.398	6.349
AI_apply	27489	1.056	1.330	0.000	0.693	6.497
lnage	27483	1.804	0.924	0.000	1.946	3.526
size	27479	0.598	0.774	-1.426	0.625	2.413
cashflow	27489	0.051	0.066	-0.141	0.049	0.247
lev	27489	0.384	0.192	0.049	0.375	0.866
asset	27484	0.335	0.156	0.030	0.324	0.742
growth	27488	0.214	0.382	-0.242	0.102	2.273
roa	27489	0.044	0.061	-0.217	0.044	0.210

3.3 Specification of the Econometric Framework

The empirical model relates firm-level innovation to the intensity of AI adoption as follows: $$Innov_{it} = \alpha_0 + \alpha_1 AI_apply_{it} + \gamma X_{it} + \lambda_i + \mu_t + \varepsilon_{it} \quad (1)$$

The subscripts *i* and *t* denote firms and years, respectively; $Innov_{it}$ represents the innovation output of firm; AI_apply_{it} measures the intensity of AI technology adoption at the firm level; X_{it} is a vector of control variables; λ_i is the firm fixed effect; while year effects, μ_t, account for shocks common to all firms in a given year. which absorbs common time trends; The residual is denoted by ε_{it}. Statistical inference allows for correlated errors within industry-year cells.

4. Research Findings

4.1 Baseline Evidence

Table 2 introduces the explanatory variables and fixed effects in stages. Column (1) includes neither fixed effects nor covariates, with inference based on industry-year clusters. The point estimate on AI adoption is 0.133 and differs from zero at the 1% level. This estimate captures the unconditional positive association between AI use and firms'overall innovation output.

Column (2) absorbs firm and year effects. The AI coefficient falls to 0.022 but remains statistically different from zero at the 5% level. This attenuation suggests that persistent firm heterogeneity and common temporal shocks explain a considerable portion of the raw association. Nevertheless, the conditional relationship between AI adoption and innovation remains positive.

Column (3) adds the complete set of covariates. The coefficient remains 0.022

and retains significance at the 5% level. Its stability relative to column (2) indicates that the observed relationship is not sensitive to the inclusion of these controls.

Evaluated at the sample standard deviation of AI adoption, the coefficient in column (3) corresponds to a 0.022 increase in overall innovation output, or 1.1% of its sample mean. Given how limited the diffusion of AI remains among firms at this stage, the effect carries practical economic weight. As for model fit, the adjusted R^2 rises steadily from 0.021 to 0.683, which shows that the fixed effects have considerable explanatory power.

Taken together, the evidence supports H1: AI adoption raises firm innovation output. This relationship persists after absorbing firm and year effects and conditioning on the full set of covariates.

Table 2 Baseline Results

	(1) ln_total	(2) ln_total	(3) ln_total
AI_apply	0.133***	0.022**	0.022**
	(0.020)	(0.009)	(0.009)
lnage			-0.029*
			(0.017)
size			-0.032**
			(0.016)
cashflow			-0.285***
			(0.091)
lev			0.200***
			(0.055)
asset			-0.073
			(0.066)
growth			-0.026
			(0.016)
roa			0.901***
			(0.117)
Year FE	—	√	√
Firm FE	—	√	√
N	27489	27489	27473
Adj. R2	0.021	0.682	0.683

Standard errors, clustered at the industry-year level, are reported in parentheses. ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively; the same convention is used in all subsequent tables.

4.2 Results by Innovation Category

Table 3 separates the aggregate response by patent class and knowledge distance. For invention patents (ln_breakthrough), the point estimate on AI use is 0.039 and is statistically different from zero. By contrast, the estimate for utility-model patents (ln_incremental) is –0.004 and cannot be distinguished statistically from zero. AI thus clearly raises breakthrough innovation without encouraging low-end incremental innovation. One explanation is that invention patents face stricter examination and a higher inventive bar, which aligns with AI's strengths in combining knowledge and recognizing patterns. Utility-model patents, by contrast, embody relatively little technology, and improving them depends more on accumulated engineering experience, leaving limited room for AI to contribute.

Dividing patents by their distance from a firm’s existing knowledge base

produces the same pattern. For innovation outside established domains, the point estimate is 0.008 (ln_explore) and cannot be distinguished from zero. Innovation within familiar domains responds more strongly, with an estimated effect of 0.021 (ln_exploit) and $p < 0.05$. AI thus reinforces what firms already do in familiar technological areas, but it has yet to push them into unfamiliar ones. This asymmetry follows from how AI works. AI learns from historical data and optimizes on that basis, so it does well in known fields where data are plentiful and patterns are clear. In new fields, where there is no prior data to learn from, its forecasting and decision-making run into hard limits.

Putting the four sets of results together, a clear pattern emerges: AI helps firms achieve high-quality breakthroughs along the technological trajectory they are already on—a model of focus, deepening, and quality upgrading. The results support one another, and they show that AI does not lead firms to diversify blindly into new fields. Instead, it helps them go deeper and further along the paths where they already have expertise.

Table 3 Effects by Innovation Type

	(1) ln_breakthrough	(2) ln_incremental	(3) ln_explore	(4) ln_exploit
AI_apply	0.039***	-0.004	0.008	0.021**
	(0.010)	(0.009)	(0.010)	(0.010)
lnage	-0.053***	0.008	-0.014	-0.019
	(0.018)	(0.018)	(0.015)	(0.017)
size	-0.028*	-0.032*	-0.008	-0.039**
	(0.016)	(0.019)	(0.016)	(0.017)
cashflow	-0.292***	-0.108	-0.234**	-0.243***
	(0.098)	(0.103)	(0.095)	(0.092)
lev	0.236***	0.169***	0.174***	0.181***
	(0.058)	(0.060)	(0.058)	(0.054)
asset	-0.034	-0.049	0.013	-0.058
	(0.068)	(0.073)	(0.066)	(0.066)
growth	-0.025	-0.015	0.018	-0.026*
	(0.015)	(0.018)	(0.016)	(0.016)
roa	1.002***	0.609***	0.690***	0.866***
	(0.124)	(0.135)	(0.118)	(0.120)
Year FE	√	√	√	√
Firm FE	√	√	√	√
N	27473	27473	27473	27473
Adj. R2	0.662	0.692	0.373	0.695

4.3 Addressing Endogeneity

The baseline model absorbs both dimensions of fixed effects and conditions on numerous firm attributes, yet endogeneity remains possible. On the one hand, enterprises with high innovation output usually have stronger technology absorption capacity and more abundant innovation resources, and may consequently adopt AI technology more intensively, creating reverse-causality concerns. On the other hand, despite the fact that this paper has controlled multiple observable variables, unobservable factors such as managers' technological preferences may simultaneously affect AI application and innovation output, giving rise to omitted-variable bias. To confront these concerns and sharpen the causal

interpretation, the analysis turns to an instrumental-variable strategy.

Consistent with prior work, a Bartik-style instrument is built by interacting the initial industry-level intensity of AI application with national growth in the tertiary sector. Specifically, the industry's mean level of AI application in 2009 provides the initial exposure share, while average growth in tertiary-sector value added across all other cities supplies the aggregate shock; the product of the two yields a firm-year instrument. The relevance condition is plausible because changes in tertiary-sector value added reflect the dynamism of the service economy, which shapes firms' operating environment and market demand, thereby affecting their adoption of new technologies, AI in particular. The exclusion restriction is plausible, since shifts in tertiary-sector value added elsewhere are unlikely to influence local firms' innovation output directly.

Instrumental-variable estimates appear in Table 4. In the first-stage equation, the excluded instrument enters with an estimate of -5.014 ($p < 0.01$), indicating substantial predictive power for AI adoption. The Kleibergen-Paap rk LM statistic equals 164.36 ($p < 0.001$), allowing the null of underidentification to be rejected. The Cragg-Donald Wald F statistic reaches 141.94; by comparison, the relevant Stock-Yogo threshold is 16.38 under the 10% maximal IV size criterion. Weak-instrument concerns are therefore unlikely to account for the results. Turning to the second stage, the estimated response of innovation to AI adoption is 0.331 ($p < 0.01$). Its positive sign and statistical precision are qualitatively consistent with the baseline evidence, reinforcing the conclusion that AI adoption promotes firm innovation.

Table 4 Instrumental-Variable Estimates

	(1) AI_apply_A	(2) ln_total
IV_Bartik_ind	-5.014***	
	(1.212)	
AI_apply_A		0.331***
		(0.102)
lnage	0.074***	-0.060***
	(0.016)	(0.017)
size	-0.015	-0.023
	(0.015)	(0.014)
cashflow	-0.075	-0.286***
	(0.099)	(0.098)
lev	0.332***	0.110*
	(0.055)	(0.067)
asset	-0.412***	0.071
	(0.056)	(0.077)
growth	-0.039***	-0.013
	(0.015)	(0.016)
roa	-0.243**	1.052***
	(0.119)	(0.122)
Year FE	√	√
Firm FE	√	√
N	25344	25344
Adj. R2	0.810	-0.237

Kleibergen-Paap rk LM = 164.36 ($p<0.001$); Cragg-Donald Wald F = 141.94, well above the 10% Stock-Yogo critical value.

4.4 Robustness Tests

The stability of the baseline result is assessed along six dimensions: variable measurement, model setting, standard misinference, missing variables, sample interval, and extreme values.

Column (1) replaces the baseline AI measure with AI_apply_B, an index derived exclusively from the Management Discussion and Analysis (MD&A) section. This section records management's assessment of current operations and prospective strategy and therefore provides a more focused signal of the firm's intention to adopt AI than the full annual report. The resulting estimate is 0.019 ($p < 0.05$), indicating that the main finding is not tied to the original text-based measure.

Column (2) absorbs firm-by-industry effects, thereby removing unobserved attributes that remain constant within firm-industry pairs. The estimate for AI_apply is 0.022 ($p < 0.05$), identical to the baseline estimate. Time-invariant omitted factors at this level are therefore unlikely to explain the result.

Column (3) recalculates inference using firm-level clusters. The point estimate remains 0.022, although its p-value rises above 0.05 but remains below 0.10. The modest loss of precision suggests that the baseline significance is not an artifact of clustering errors by industry-year.

Column (4) augments the specification with city-level GDP, GDP growth, science expenditure, and education expenditure. Under this specification, AI_apply carries an estimate of 0.026 ($p < 0.01$). The additional city characteristics are individually insignificant, and their inclusion leaves the main result intact.

Column (5) removes observations from 2019 onward to avoid contamination from the AI pilot-zone policy. The estimate for AI_apply_A rises to 0.045 ($p < 0.01$), providing no indication that the policy drives the baseline finding.

Column (6) trims the AI word-frequency measure at its 1st and 99th percentiles. This adjustment produces an estimate of 0.021 ($p < 0.05$), showing little sensitivity to unusually large or small text-based values.

Across these exercises, the point estimates remain positive, statistically distinguishable from zero at conventional levels, and confined to a relatively narrow range of 0.019-0.045. Their proximity to the baseline estimate of 0.022 indicates that the main result survives alternative measurement, specification, inference, sample, and outlier treatments.

Table 5 Robustness Checks

	(1) ln_total	(2) ln_total	(3) ln_total	(4) ln_total	(5) ln_total	(6) ln_total
AI_apply_B	0.019**					
	(0.009)					
AI_apply		0.022**	0.022*	0.026***	0.045***	
		(0.009)	(0.012)	(0.009)	(0.013)	
AI_apply_w						0.021**
						(0.009)
lnage	-0.033**	-0.029*	-0.029	-0.029*	-0.031	-0.029*
	(0.017)	(0.017)	(0.020)	(0.017)	(0.027)	(0.017)
size	-0.031*	-0.032**	-0.032*	-0.035**	-0.049**	-0.032**
	(0.017)	(0.016)	(0.019)	(0.017)	(0.021)	(0.016)
cashflow	-0.271***	-0.285***	-0.285***	-0.290***	-0.357**	-0.285***
	(0.091)	(0.091)	(0.100)	(0.092)	(0.148)	(0.091)
lev	0.210***	0.200***	0.200**	0.202***	0.159*	0.200***
	(0.055)	(0.055)	(0.084)	(0.055)	(0.083)	(0.055)

asset	-0.082	-0.073	-0.073	-0.079	-0.001	-0.073
	(0.066)	(0.066)	(0.093)	(0.066)	(0.093)	(0.066)
growth	-0.026	-0.026	-0.026*	-0.029*	-0.035*	-0.026
	(0.016)	(0.016)	(0.015)	(0.016)	(0.021)	(0.016)
roa	0.912***	0.901***	0.901***	0.895***	0.947***	0.901***
	(0.117)	(0.117)	(0.136)	(0.117)	(0.182)	(0.117)
gdp				-0.000		
				(0.000)		
gdp_growth				0.005		
				(0.003)		
science_exp				0.000		
				(0.000)		
edu_exp				-0.000		
				(0.000)		
Year FE	√	√	√	√	√	√
Firm FE	√	√	√	√	√	√
Firm-Industry FE	—	√	—	—	—	—
N	27244	27473	27473	27290	12572	27473
Adj. R2	0.683	0.683	0.683	0.683	0.707	0.683

Column (1): AI_apply_A is replaced by AI_apply_B. Column (2): firm-industry fixed effects added. Column (3): standard errors clustered at the firm level. Column (4): city-level controls added (GDP, GDP growth, science and education spending). Column (5): observations from 2019 onwards excluded. Column (6): AI_apply_A winsorized at 1% and 99%. Except in Column (3),

4.5 Heterogeneity Analysis

1. Enterprise scale. Enterprise scale reflects the difference between resource endowment and technology absorption capacity. Large-scale enterprises have more abundant funds, more perfect data infrastructure, and more mature technical teams, which can bear the fixed costs and trial-and-error risks of AI deployment. Smaller firms, by contrast, tend to operate under tighter resource constraints. Even if AI is introduced, it may be challenging to produce significant innovation due to insufficient complementary assets. Firms are classified by comparing the logarithm of operating revenue with the corresponding industry-year median. The estimates reported in columns (1) and (2) of Table 6 display a clear size-related pattern. For larger firms, the AI coefficient is 0.038 ($p < 0.01$), whereas the estimate for smaller firms is -0.002 and indistinguishable from zero. These results suggest that the innovation gains associated with AI accrue primarily to firms with sufficient complementary resources.

2. Density of industrial elements. Significant differences exist in the technical foundation, research and development investment, and human capital structure of industries with various element intensities, which may affect the efficiency of enterprises' absorption and transformation of AI technology. Technology-intensive sectors invest more intensively in R&D, have a richer high-skilled labor force, and possess stronger technological innovation capabilities, which can better complement AI technology so as to transform technological potential into innovative achievements. Drawing on the classification methods of existing research, the sample industries are grouped as labor-intensive, capital-intensive, or technology-intensive. Columns (3)–(5) of Table 6 show that the estimate for the technology-intensive group is 0.049 ($p < 0.01$). By comparison, the estimates obtained for the labor-and-capital-intensive groups cannot be distinguished from zero at conventional significance levels. The innovation gains from AI are thus concentrated in technology-intensive industries, where the existing technological base enables AI to deliver its effects.

3. Industry competitiveness. The competitive environment in the market is an essential factor affecting the technology adoption and innovation incentives of enterprises. According to the theory of industrial organization, competitive pressure will prompt enterprises to seek technological progress for competitive advantages; while in a market lacking competition, enterprises lack the urgency to introduce new technologies to improve innovative capability. The sample is divided into high-and-low-competition groups using the industry-year median of the Herfindahl index. As reported in columns (6) and (7) of Table 6, the estimate for firms operating in more competitive industries is 0.040 ($p < 0.01$). The corresponding estimate in less competitive industries is 0.008 and statistically indistinguishable from zero. This pattern is consistent with competitive pressure strengthening firms' incentives to translate AI adoption into innovative output.

Taken together, the subgroup estimates associate AI with greater innovation primarily among larger firms, technology-intensive industries, and more competitive product markets. The pattern indicates that the innovative returns to this general-purpose technology depend on complementary firm resources, sectoral technological capabilities, and external competitive incentives. The results thus clarify the settings in which AI adoption is most likely to generate measurable innovation gains.

Table 6 Heterogeneous Effects across Firms and Industries

	(1) ln_total	(2) ln_total	(3) ln_total	(4) ln_total	(5) ln_total	(6) ln_total	(7) ln_total
AI_apply	0.038***	-0.002	0.033	-0.022	0.049***	0.040***	0.008
	(0.013)	(0.015)	(0.024)	(0.016)	(0.013)	(0.014)	(0.014)
lnage	-0.016	-0.076**	-0.038	-0.012	-0.040	0.002	-0.054**
	(0.023)	(0.033)	(0.040)	(0.025)	(0.027)	(0.024)	(0.023)
size	-0.037	-0.039	-0.051	-0.014	-0.058**	0.013	-0.070***
	(0.025)	(0.032)	(0.039)	(0.023)	(0.027)	(0.025)	(0.022)
cashflow	-0.330***	-0.247*	-0.372	-0.471***	-0.028	-0.229**	-0.344**
	(0.123)	(0.150)	(0.275)	(0.136)	(0.131)	(0.115)	(0.135)
lev	0.074	0.377***	0.144	0.116	0.190**	0.041	0.267***
	(0.082)	(0.109)	(0.153)	(0.073)	(0.092)	(0.076)	(0.085)
asset	-0.072	-0.125	-0.036	-0.198**	0.046	-0.113	-0.111
	(0.095)	(0.105)	(0.182)	(0.091)	(0.109)	(0.098)	(0.094)
growth	-0.029	-0.012	0.005	-0.012	-0.032	-0.018	-0.035
	(0.019)	(0.033)	(0.050)	(0.021)	(0.023)	(0.021)	(0.022)
roa	0.725***	0.953***	0.931***	0.932***	0.814***	0.754***	0.843***
	(0.168)	(0.190)	(0.282)	(0.186)	(0.174)	(0.166)	(0.167)
_cons	2.572***	2.836***	2.399***	2.629***	2.798***	2.644***	2.725***
	(0.054)	(0.081)	(0.111)	(0.055)	(0.060)	(0.053)	(0.060)
Year FE	√	√	√	√	√	√	√
Firm FE	√	√	√	√	√	√	√
N	13203	13192	3623	11837	11677	13645	13301
Adj. R2	0.670	0.710	0.668	0.661	0.706	0.712	0.690

The table presents regressions estimated on subsamples. Columns (1)-(2) split the sample by firm size, Columns (3)-(5) by industry factor intensity, and Columns (6)-(7) by industry competition.

5. Mechanism Analysis

Table 7 reports estimates for six mechanism measures. AI does not directly promote enterprise innovation, but works by changing resource allocation and knowledge flow within the enterprise. The analysis traces these channels across the input, process, and allocation margins.

5.1 R&D Investment and Human Capital Upgrading

To fully play the role of AI, enterprises need to make adjustments in talent and capital. On the labor side, developing and deploying AI systems require people who understand algorithms, data science, and software engineering. This will encourage enterprises to recruit more R&D personnel and make the talent structure more inclined to high skills. On the capital side, adopting AI requires enterprises to build data systems, training models, and develop applications, which require additional R&D investment. As AI is knowledge-intensive, firms that wish to realize its potential typically raise their research spending. The paper therefore uses the proportions of R&D personnel and R&D spending to capture changes on the input side.

For R&D investment intensity, column (1) reports an AI estimate of 0.696, different from zero at the 10% threshold, which indicates greater AI use raises corporate R&D expenditure. Column (2) yields an AI estimate for R&D personnel of 0.303, with $p < 0.01$, indicating that AI has significantly improved the talent structure of enterprises. Overall, the main effect of AI on the input side is to upgrade the talent structure of the enterprise, rather than to significantly expand R&D investment. The possible reason is that AI can be realized by buying ready-made products, without investment in numerous R&D. However, high-quality R&D talents are necessary conditions to adapt to the intelligent transformation.

5.2 Knowledge Spillover and Knowledge Diversity

As an information-processing tool, AI can change how knowledge moves and combines between firms. In the past, the spread of knowledge was limited by individual cognitive capacity and by information asymmetry, and it largely worked as a one-way, point-to-point transfer. AI uses algorithms to identify and match information, which turns knowledge diffusion into a networked, shared process and so strengthens its diffusion effect. Different innovators can then identify, share and update knowledge on the same platform. More specifically, natural language processing and knowledge-mapping techniques can pull meaning out of cluttered information and link it together in a structured way, helping firms find outside knowledge and partners. AI can also help firms combine knowledge more richly within technologies they already have and thereby widen their knowledge base. Knowledge spillovers are measured with patent citations: the more citations a firm's patents receive, the wider its knowledge spillovers and the greater its influence in the industry. Knowledge diversity is captured using how many patent applications a firm files within familiar IPC subclasses, which reflects how varied its knowledge combinations are inside the technological boundaries it already occupies.

Column (3) reports an AI estimate for high-quality innovation of 0.024, with $p < 0.01$, which indicates greater AI use increases the production of highly cited patents and supports the creation and diffusion of high-quality knowledge, confirming the knowledge-spillover channel. Column (4) reports a coefficient of 0.026 on patent applications in familiar IPC subclasses; this differs from zero at $p < 0.01$, suggesting that AI fosters knowledge accumulation and diversification within firms' existing technology fields and thereby verifies the knowledge diversity effect.

5.3 Resource Allocation Efficiency

AI affects resource allocation in two ways. First, it improves capital allocation. Because AI can forecast, firms can judge the return on an investment project more accurately and reduce the mismatch between capital and projects. AI also connects

departments within a firm, encourages cooperation and information sharing across them, cuts duplicated labor and wasted resources, and frees more resources for innovation. Second, it changes the asset structure. AI is driven by software and algorithms and depends more on computing power, data, and talent than on plant and equipment. That encourages firms to rely less on traditional fixed assets and to shift resources towards intangibles such as software, algorithms, and data, giving them a lighter asset structure. Capital allocation is measured with the investment expenditure rate, and the fixed-asset ratio is used to capture changes in the asset structure.

Column (5) reports a coefficient of 0.001 on AI for the investment expenditure rate; it differs from zero at $p < 0.01$, suggesting that greater AI use improves capital expenditure and allocation efficiency. Column (6) places the AI estimate for the fixed-assets ratio at -0.001 ($p < 0.01$), implying that greater AI use lowers the share of fixed assets in total assets and encourages firms to move toward a lighter asset structure.

Table 7 Tests of the Mechanisms

	(1) rd_ratio	(2) rd_person	(3) cit90	(4) new_apply	(5) invt	(6) fixedasset
AI_apply	0.696*	0.303***	0.024***	0.026***	0.001***	-0.001***
	(0.404)	(0.083)	(0.006)	(0.010)	(0.000)	(0.000)
lnage	-2.256	0.062	0.011	-0.019	-0.002*	0.000
	(1.686)	(0.129)	(0.009)	(0.017)	(0.001)	(0.001)
size	1.383	0.299**	0.004	-0.039**	0.005***	-0.005***
	(2.315)	(0.129)	(0.008)	(0.017)	(0.001)	(0.001)
cashflow	-16.462	-1.298*	-0.046	-0.249***	-0.015***	0.098***
	(14.828)	(0.706)	(0.054)	(0.094)	(0.005)	(0.007)
lev	-34.802	-2.541***	0.044	0.181***	0.032***	-0.057***
	(22.345)	(0.555)	(0.028)	(0.055)	(0.004)	(0.004)
asset	11.543	-3.487***	0.012	-0.068	-0.051***	0.695***
	(12.561)	(0.609)	(0.030)	(0.067)	(0.005)	(0.009)
growth	16.796	-0.693***	0.000	-0.028*	-0.006***	-0.004***
	(11.420)	(0.136)	(0.007)	(0.016)	(0.001)	(0.001)
roa	-116.611	-1.461	0.130*	0.865***	0.072***	-0.128***
	(73.116)	(0.989)	(0.069)	(0.121)	(0.006)	(0.009)
_cons	22.584**	19.970***	0.042*	2.474***	0.058***	0.001
	(9.175)	(0.356)	(0.023)	(0.040)	(0.003)	(0.004)
Year FE	√	√	√	√	√	√
Firm FE	√	√	√	√	√	√
N	22146	17769	27473	27473	27467	27473
Adj. R2	0.263	0.908	0.358	0.704	0.461	0.938

The table reports how AI use affects the mechanism variables. Columns (1)-(2): R&D input and human capital upgrading; Columns (3)-(4): knowledge spillovers and knowledge diversity; Columns (5)-(6): resource allocation efficiency.

6. Concluding Remarks and Policy Implications

6.1 Main Findings

Drawing on a sample of publicly traded Chinese manufacturing enterprises covering 2010-2024, the study builds an enterprise-specific indicator of AI use from references to AI appearing in annual reports and examines innovation along two dimensions, patent type and knowledge distance. The analysis centers on whether AI affects exploitative and exploratory innovation differently, and on the channels

through which any such effects operate.

The central results are summarized below.

First, AI adoption raises firms' innovation output. The baseline estimates show that greater AI use is associated with more innovation at the 5% level once firm and year effects and the remaining firm-level factors are controlled for. This finding holds up under numerous additional checks using alternative measures of AI, richer fixed effects, different clustering methods, additional city-level controls, and various treatments of extreme values. The instrumental-variable estimates likewise support a causal reading.

Second, AI does not affect all types of innovation equally. The classification results show that AI adoption increases invention patents and exploitative innovation, whereas its effects on utility-model patents and exploratory innovation are statistically insignificant. These results suggest that AI mainly helps firms improve innovation that builds on knowledge and technologies they already possess.

Third, the mechanism analysis identifies five channels through which AI promotes firm innovation: improving the composition of human capital, facilitating high-quality knowledge spillovers, expanding knowledge diversity, improving capital allocation, and shifting firms toward a lighter asset structure. Specifically, the share of R&D personnel, patent citations, patent applications in familiar technological fields, the investment expenditure rate, and the fixed-assets ratio are all statistically significant. These estimates indicate that AI affects innovation through multiple pathways operating on the input, process, and allocation sides.

Fourth, the effects are stronger in some settings than in others. Large-scale firms gain more than small ones, technology-intensive industries gain more than other industries, and competitive markets amplify the effect. Scale supplies resources, technology intensity provides capability, and competition fosters incentives. These conditions collectively enable AI-driven innovation.

6.2 Policy Implications

Several policy recommendations are provided from these findings.

First, governments should keep pushing AI and corporate innovation closer together. This paper finds that AI raises innovation output, especially for invention patents and exploitative innovation. Governments should therefore improve the institutional environment for AI development, increase support for firms' use of AI, and lower the up-front cost of intelligent transformation through fiscal subsidies, tax relief and special funds. At the same time, support policies should be tailored to the differing needs of different types of firms rather than applied uniformly to all.

Second, human capital development should be strengthened to unleash the innovation potential of AI. The mechanism analysis shows that AI improves innovation by upgrading the human capital structure, facilitating knowledge spillovers, and increasing the efficiency of resource allocation. The government should strengthen the recruitment and training of professionals in AI-related fields, consolidate university-industry-research collaboration, and encourage universities, research institutions, and firms to jointly cultivate high-quality talent. Firms should proactively adjust their talent structure, increase the allocation of R&D personnel, enhance employees' literacy in AI technology, and build a human capital system adapted to intelligent transformation.

Third, policy should recognize the catalytic role that a competitive environment plays when shaping innovation effects of AI. The subgroup estimates indicate that the innovation returns associated with AI tend to be stronger among highly competitive

industries. The government should improve the market system of fair competition, break down industry monopolies and administrative barriers, and create a market environment conducive to firms' technological innovation. Furthermore, intellectual property protection should be strengthened to ensure that firms obtain reasonable returns from innovation activities and to form a long-term mechanism that incentivizes innovation.

Fourth, smaller firms need direct public support. At present, large corporations capture almost all of the innovation gains from AI, and smaller firms clearly lag behind. Governments can close this gap by setting up shared public service platforms that offer cheap computing power, technical advice and straightforward setup guides, which would lower the entry costs that normally put smaller firms off.

References

Ali, D., & Frimpong, S. (2020). Artificial intelligence, machine learning a nd process automation: Existing knowledge frontier and way forward for minin gsector. *Artificial Intelligence Review*, *53*(8), 6025~6042. https://doi.org/10.1007/s10462-020-09841-6

Audretsch, D., Belitski, M., & Fiedler, A. (2025). The knowledge spillover theory of entrepreneurship and innovation: Taking stock and new directions. *The Journal of Technology Transfer*, *51*(4), 2117~2140. https://doi.org/10.1007/s10961-025-10215-9

Benner, M. J., & Tushman, M. L. (2003). Exploitation, Exploration, and Process Management: The Productivity Dilemma Revisited. *Academy of Management Review*, *28*(2), 238~256. https://doi.org/10.5465/amr.2003.9416096

Brem, A., Giones, F., & Werle, M. (2023). The AI Digital Revolution in Innovation: A Conceptual Framework of Artificial Intelligence Technologies for the Management of Innovation. *IEEE Transactions on Engineering Management*, *70*(2), 770~776. https://doi.org/10.1109/TEM.2021.3109983

Che, S., Liu, C., Wang, J., & Wang, J. (2026). Can artificial intelligence drive enterprise green management innovation? A new perspective on harnessing intelligence. *Technology in Society*, *85*, 103156. https://doi.org/10.1016/j.techsoc.2025.103156

Chin, T., Ghouri, M. W. A., Jin, J., & Deveci, M. (2024). AI technologies affording the orchestration of ecosystem-based business models: The moderating role of AI knowledge spillover. *Humanities and Social Sciences Communications*, *11*(1), 496. https://doi.org/10.1057/s41599-024-03003-7

Choi, S., Kang, H., Kim, N., & Kim, J. (2025). The dual edges of AI: Advancing knowledge while reducing diversity. *PNAS Nexus*, *4*(5), pgaf138. https://doi.org/10.1093/pnasnexus/pgaf138

D'Amico, E., Belitski, M., Braga, A., & Savoie, R. (2025). Enhancing knowledge spillover of innovation through artificial intelligence: An empirical investigation. *The Journal of Technology Transfer*, *51*(4), 2222~2250. https://doi.org/10.1007/s10961-025-10217-7

Fonteneau, F., Mollins, J., Marchi, S., Russo, L., Gentaz, A., Daoud, M., & André, A.-A. (2025). Advancing the measurement of investments in artificial intelligence. *OECD*. https://www.oecd.org/en/publications/advancing-the-measurement-of-investments-in-artificial-intelligence_13e0da2f-en.html

Fu, Y., Ni, J., & Fang, M. (2025). The impact of artificial intelligence on digital enterprise innovation. *Journal of Strategy & Innovation*, *36*(1), 200538. https://doi.org/10.1016/j.jsinno.2025.200538

Gao, Y., Liu, S., & Yang, L. (2025). Artificial intelligence and innovation capability: A dynamic capabilities perspective. *International Review of Economics & Finance*, *98*, 103923. https://doi.org/10.1016/j.iref.2025.103923

Ghimire, A. (2025). Utilizing ChatGPT to integrate world English and diverse knowledge: A transnational perspective in critical artificial intelligence (AI) literacy. *Computers and Composition*, *75*, 102913. https://doi.org/10.1016/j.compcom.2024.102913

Grashof, N., & Kopka, A. (2022). Artificial intelligence and radical innovation: An opportunity for all companies? *Small Business Economics*, *61*(2), 771~797.

https://doi.org/10.1007/s11187-022-00698-3

Guan, C., Mou, J., & Jiang, Z. (2020). Artificial intelligence innovation in education: A twenty-year data-driven historical analysis. *International Journal of Innovation Studies*, *4*(4), 134~147. https://doi.org/10.1016/j.ijis.2020.09.001

Haefner, N., Wincent, J., Parida, V., & Gassmann, O. (2021). Artificial intelligence and innovation management: A review, framework, and research agenda☆. *Technological Forecasting and Social Change*, *162*, 120392. https://doi.org/10.1016/j.techfore.2020.120392

Han, F., & Mao, X. (2024). Artificial intelligence empowers enterprise innovation: Evidence from China's industrial enterprises. *Applied Economics*, *56*(57), 7971~7986. https://doi.org/10.1080/00036846.2023.2289916

Hassannataj Joloudari, J., Mojrian, S., Saadatfar, H., Nodehi, I., Fazl, F., Khanjani Shirkharkolaie, S., Alizadehsani, R., Kabir, H. M. D., Tan, R.-S., & Acharya, U. R. (2024). Resource allocation problem and artificial intelligence: The state-of-the-art review (2009–2023) and open research challenges. *Multimedia Tools and Applications*, *83*(26), 67953 ~ 67996. https://doi.org/10.1007/s11042-024-18123-0

Hutchinson, P. (2021). Reinventing Innovation Management: The Impact of Self-Innovating Artificial Intelligence. *IEEE Transactions on Engineering Management*, *68*(2), 628~639. https://doi.org/10.1109/TEM.2020.2977222

Jiang, L., Xuan, Y., & Zhang, K. (2024). Unlocking innovation potential: The impact of artificial intelligence transformation on enterprise innovation capacity. *European Journal of Innovation Management*, *28*(8), 4112~4131. https://doi.org/10.1108/EJIM-07-2024-0809

Li, B., Wang, L., Fu, S., & Wan, J. (2025). How does artificial intelligence empower enterprise technological innovation: From the perspective of knowledge management capability. *Journal of Knowledge Management*, 1~31. https://doi.org/10.1108/JKM-03-2025-0345

Li, D., Xiao, J., & Yang, F. (2024). Artificial Intelligence and Enterprise Green Innovation: Intrinsic Mechanisms and Heterogeneous Effects. *Sustainability*, *16*(21), 9246. https://doi.org/10.3390/su16219246

Liang, G., Liang, Y., Niu, D., & Shaheen, M. (2024). Balancing sustainability and innovation: The role of artificial intelligence in shaping mining practices for sustainable mining development. *Resources Policy*, *90*, 104793. https://doi.org/10.1016/j.resourpol.2024.104793

Liu, J., Chang, H., Forrest, J. Y.-L., & Yang, B. (2020). Influence of artificial intelligence on technological innovation: Evidence from the panel data of china's manufacturing sectors. *Technological Forecasting and Social Change*, *158*, 120142. https://doi.org/10.1016/j.techfore.2020.120142

Lu, H., Du, D., & Qin, X. (2022). Assessing the Dual Innovation Capability of National Innovation System: Empirical Evidence from 65 Countries. *Systems*, *10*(2), 23. https://doi.org/10.3390/systems10020023

Lu, Q., Qin, W., & Yang, Y. (2026). Intelligent resource orchestration counselor: How artificial intelligence enables enterprises' exploratory innovation. *Journal of Business & Industrial Marketing*, *41*(3), 396~412. https://doi.org/10.1108/JBIM-04-2025-0318

Ma, J., Shang, Y., & Liang, Z. (2025). Digital transformation, artificial intelligence and enterprise innovation performance. *Finance Research Letters*, *78*, 107190. https://doi.org/10.1016/j.frl.2025.107190

March, J. G. (1991). Exploration and Exploitation in Organizational Learning. *Organization Science*, *2*(1), 71~87. https://doi.org/10.1287/orsc.2.1.71

Mühlroth, C., & Grottke, M. (2022). Artificial Intelligence in Innovation: How to Spot Emerging Trends and Technologies. *IEEE Transactions on Engineering Management*, *69*(2), 493~510. https://doi.org/10.1109/TEM.2020.2989214

Nagarajan, G. (2022). Optimizing Project Resource Allocation through a Caching-Enhanced Cloud AI Decision Support System. *International Journal of Computer Technology and Electronics Communication*, *5*(2), 4812~4820. https://doi.org/10.15680/IJCTECE.2022.0502003

Paschen, U., Pitt, C., & Kietzmann, J. (2020). Artificial intelligence: Building blocks and an innovation typology. *Business Horizons, ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING*, *63*(2), 147~155. https://doi.org/10.1016/j.bushor. 2019.10.004

Philip, P. G. (2024). Evaluating the Impact of AI-Powered Resource Allocation Systems on Project Efficiency and Cost Optimization. *The USA Journals TAJET*, *6*(03), 31~44.

Sjödin, D., Parida, V., & Kohtamäki, M. (2023). Artificial intelligence enabling circular business model innovation in digital servitization: Conceptualizing dynamic capabilities, AI capacities, business models and effects. *Technological Forecasting and Social Change*, *197*, 122903. https://doi.org/10.1016/j.techfore.2023.122903

Stinson, C., & Vlaad, S. (2024). A feeling for the algorithm: Diversity, expertise, and artificial intelligence. *Big Data & Society*, *11*(1), 20539517231224247. https://doi.org/10.1177/20539517231224247

Tang, X., Li, X., Ding, Y., Song, M., & Bu, Y. (2020). The pace of artificial intelligence innovations: Speed, talent, and trial-and-error. *Journal of Informetrics*, *14*(4), 101094. https://doi.org/10.1016/j.joi.2020.101094

Verganti, R., Vendraminelli, L., & Iansiti, M. (2020). Innovation and Design in the Age of Artificial Intelligence. *Journal of Product Innovation Management*, *37*(3), 212~227. https://doi.org/10.1111/jpim.12523

Wang, J., Wang, A., Luo, K., & Nie, Y. (2024). Can artificial intelligence improve enterprise environmental performance: Evidence from China. *Journal of Environmental Management*, *370*, 123079. https://doi.org/10.1016/j.jenvman. 2024.123079

Xiao, L., Xiao, Y., & He, R. (2024). The Research and Development Investment Management in Technology Enterprises Under Artificial Intelligence. *IEEE Access*, *12*, 190936~190944. https://doi.org/10.1109/ACCESS.2024.3518070

Zhang, L., Li, H., Zhou, W., Qiu, H., & Wu, Y. J. (2025). Exploratory and Exploitative Innovation Performance in the Artificial Intelligence Industry in China from the Perspective of a Collaboration Network: A Data-Driven Analysis. *Entropy*, *27*(6), 577. https://doi.org/10.3390/e27060577

Zhang, Q., & Li, Y. (2026). Artificial intelligence and corporate innovation in small and medium-sized enterprises. *Technology in Society*, *86*, 103320. https://doi.org/10.1016/j.techsoc.2026.103320

Zhang, Q., Zuo, J., & Yang, S. (2025). Research on the impact of generative artificial intelligence (GenAI) on enterprise innovation performance: A knowledge management perspective. *Journal of Knowledge Management*, *29*(7), 2238~2257. https://doi.org/10.1108/JKM-10-2024-1198

Zhang, Y., Chang, Y., Zhang, X., & Zhou, H. (2025). Artificial Intelligence and Enterprise Technological Innovation: Evidence From China. *Managerial and Decision Economics*, *46*(7), 4056~4069. https://doi.org/10.1002/mde.70002

Zhao, J. (2021). Knowledge management capability and technology uncertainty: Driving factors of dual innovation. *Technology Analysis & Strategic Management*, *33*(7), 783~796. https://doi.org/10.1080/09537325.2020.1841896

Zhou, Y. (2024). Natural resources and green economic growth: A pathway to innovation and digital transformation in the mining industry. *Resources Policy*, *90*, 104667. https://doi.org/10.1016/j.resourpol.2024.104667