

AI Policy Shocks and Firm Pricing Power: A Quasi-Experimental Evaluation from the Involution Perspective

Ziqi Wang

Jiangsu Provincial Academy of Social Sciences

Qian Li

Nanjing University of Finance & Economics

Youmeng Kong

High School Affiliated to Nanjing University

Abstract: A key economic question in the recent debate on involutionary competition is why firms' efficiency improvements have not translated into higher price markups. To address this question, this paper treats China's National New-Generation Artificial Intelligence Innovation Development Pilot Zones as a quasi-natural experiment. Using panel data of A-share listed manufacturing firms from 2014 to 2024, it employs difference-in-differences and triple-difference methods to estimate the impact of AI policy shocks on firm markup and examines the moderating role of firm AI technology penetration. The results show that the establishment of AI pilot zones raises firm markup by approximately 0.55 percentage points on average, indicating that the policy has a modest alleviating effect on involutionary pressure. The triple-difference results demonstrate that the markup-enhancing effect of the AI policy strengthens as firm AI technology penetration increases; however, firms with insufficient technological accumulation benefit less, and the policy dividend may even widen gaps among firms. Heterogeneity analysis reveals that the above moderating effect is more pronounced among private firms and those located in eastern regions. Mechanism tests indicate that firms with higher AI technology penetration increase their AI invention patent applications after the policy shock. Further analysis shows that high-R&D-intensity firms experience significant efficiency gains from the AI policy, whereas low-R&D-intensity firms do not exhibit comparable improvements. These findings uncover the micro-level distributional mechanism of AI policy effects and provide empirical evidence for the design of differentiated industrial policies.

Keywords: Artificial Intelligence; Technology Penetration; Firm Markup; Triple Difference; Absorptive Capacity

JEL Classification: F062.4

Document Code: A

Article ID: 1000-1052(2026)00-0000-00

1. Introduction

The rapid penetration of artificial intelligence (AI) technologies is profoundly reshaping the global industrial competitive landscape. Breakthroughs in large language models and generative AI technologies, represented by ChatGPT and DeepSeek, have moved the automation of knowledge work from theory to reality, triggering widespread discussion on technological unemployment and industrial involution.^① However, the use of "involution" in economic discourse often remains at the level of metaphor. That firms improve efficiency yet fail to convert cost advantages into price markups is, in the final analysis, a problem of lost pricing power. The systematic research on firm markup by De Loecker et al. (2020) [25] provides analytical tools for understanding this phenomenon. Only by reducing involution to the dynamics of firm markup can we obtain a rigorous theoretical explanation within the analytical framework of industrial organization economics.

Echoing the above core question, the existing literature provides two key pathways for dialogue. One pathway originates from absorptive capacity theory. Cohen and Levinthal (1990) [24] demonstrated that a firm's ability to benefit from external technological shocks is largely constrained by its prior stock of relevant knowledge, implying that the effects of AI policy must necessarily manifest as a heterogeneous pattern conditioned on firms' technological readiness. The other pathway derives from research on changes in the structure of firm market power in the digital economy era. Autor et al. (2020) [22] revealed that digital technologies, by virtue of economies of scale and network effects, enable leading firms to continuously expand their market power through data advantages, driving up industry concentration, while the distribution of benefits from technology diffusion is far from uniform. Bringing these two research strands together on the question of the micro-level effects of AI policy yields a key research proposition: Which type of firm will capture the dividends brought by AI policy? Is it firms that already possess a technological foundation that receive an added bonus, or technologically weak firms that receive timely assistance?

The answer to the above question carries direct policy implications for the precise design of AI industrial policy. Since 2019, China has established AI Innovation Development Pilot Zones in three batches: the first batch covered Beijing and Shanghai; the second batch expanded to 12 cities including Tianjin, Shenzhen, Hangzhou, and Chengdu; the third batch included Shenyang, Harbin, and Zhengzhou. The underlying logic of this policy arrangement is to promote technology diffusion by reducing the cost of AI technology acquisition, constructing industrial ecosystems, and pooling talent resources. However, when policies are open to all firms in an undifferentiated manner, asymmetries in firms' absorptive capacity will lead to the differential allocation of policy effects. This concerns both the efficiency dimension and the equity dimension.

Preliminary empirical studies on the above topics have yielded important clues. Mao and Sun (2025) [8], using the AI Pilot Zone policy as a quasi-natural experiment, found that the establishment of pilot zones significantly improved firms' labor allocation efficiency, with AI technology penetration constituting one of the core transmission pathways. Zhao and Liu (2025) [18] introduced the concept of AI penetration rate and found a U-shaped relationship between AI penetration and firm innovation efficiency, with regional human capital levels positively moderating this relationship. Further research by Wang et al. (2026) [10] found that the AI application of core firms generates bilateral spillover effects in the supply chain, driving the markup distribution from a polarized state toward greater coordination. The above studies demonstrate that AI policy effects do objectively exist, but they are neither monotonic nor uniform in direction or group distribution. However, the existing literature has mostly focused on labor allocation efficiency and innovation efficiency, with relatively little direct examination of firms' pricing power in product markets—and pricing power is precisely the core dimension for determining whether firms can translate technological efficiency gains into actual market power.

To answer the above questions, this paper leverages the establishment of AI Innovation Development Pilot Zones as an exogenous policy shock. Based on panel data of A-share listed manufacturing firms from 2014 to 2024, it measures firm AI technology penetration using the frequency of AI-related terms in corporate annual reports, and employs the triple-difference (DDD) method to identify the causal moderating effect of AI technology penetration on changes in firm markup under the policy shock. The core advantage of the DDD strategy relative to conventional grouped DID is that it does not rely on the parallel trends assumption for the outcome variable between treatment and control groups; instead, by introducing a firm-level moderating variable, it transforms identification into a continuous interaction between treatment intensity and policy exposure, requiring only that firms at different levels of AI technology penetration exhibit parallel trends in markup changes in the counterfactual state—a much weaker assumption.

The analysis of this paper takes absorptive capacity theory as its fundamental framework. The AI Pilot Zone policy reduces the cost for firms to access external AI technological knowledge, but whether this cost advantage can be translated into pricing power depends on whether firms themselves possess the conditions for recognizing, assimilating, and utilizing external knowledge. Zahra and George (2002) [27]

distinguished absorptive capacity into two dimensions: potential absorptive capacity and realized absorptive capacity. The AI technology penetration of concern in this paper—the level of AI technology attention and knowledge reserves reflected in firms' annual reports—belongs to the category of potential absorptive capacity. Firms with higher penetration can, under the policy's influence, initiate AI technology innovation activities earlier and construct market barriers through two pathways—product differentiation and production cost optimization—thereby maintaining or enhancing their pricing power in an increasingly competitive environment. Firms with lower penetration, lacking the necessary knowledge reserves, are constrained in their absorption and utilization of policy resources, and may even face further compression of their profit margins by technologically leading firms.

The contributions of this paper are reflected at three levels. First, in terms of research perspective, it incorporates firm AI technology penetration into the AI policy effect evaluation framework, conducting analysis from the angle of micro-level heterogeneity and its formation mechanisms, thereby moving beyond the scope of average treatment effect discussions. Second, with respect to identification methods, it exploits the batched rollout of the AI Pilot Zone policy, combined with cross-sectional differences in AI technology penetration across firms, to construct a DDD identification scheme, providing a reusable analytical tool for similar research topics. Third, in terms of empirical evidence, this paper delineates a complete causal chain encompassing policy effect identification, absorptive capacity testing, and mechanism verification. The main findings include: the moderating effect of AI technology penetration is more pronounced among private firms and in eastern regions; this effect operates through the channel of increased AI invention patent applications; and R&D intensity is a key prerequisite for AI policy to generate efficiency-enhancing effects, as firms with insufficient R&D investment struggle to obtain substantive benefits from the policy.

2. Literature Review and Research Hypotheses

2.1 Literature Review

Can the economic benefits of AI technology be equally captured by all firms? This is precisely the core question this paper seeks to answer, and also the logical starting point that connects three strands of literature.

The first strand concerns the micro-level heterogeneity of AI's macro effects. Macro-level studies focus on the average treatment effects of AI on employment, economic growth, and factor income shares (Acemoglu and Restrepo, 2018; Aghion et al., 2017) [20–21], but these average effects cannot reflect the enormous heterogeneity existing at the firm level. Autor et al. (2020) [22] revealed the "superstar firm" phenomenon in the digital technology domain: leading firms, leveraging data advantages and technological barriers, continuously expand their market power and drive up industry concentration, while the profit margins of lagging firms are further squeezed. This indicates that the diffusion of AI technology is not neutral in a competitive sense, and the distribution of its benefits exhibits pronounced asymmetric characteristics. In the Chinese context, Cai and Chen (2019) [1] examined the differential employment impacts of AI on different skill groups and industries; Li et al. (2023) [6], based on manufacturing firm data, found that AI application significantly increases firm total factor productivity but that the effect magnitude varies significantly across different types of firms. Liu and Cheng (2023) [7], from the perspective of global production networks, revealed the asymmetric nature of technology spillovers from digital industries, finding that the technology spillovers provided by digital industries are significantly greater than the spillovers they absorb. This finding raises a key policy question: If AI dividends naturally favor the strong, can industrial policies aimed at promoting AI technology diffusion alter this distributional pattern, or will they exacerbate it?

Around the above question, quasi-experimental studies using China's AI Innovation Development Pilot Zones as the research object are progressively providing preliminary yet important empirical evidence. Mao and Sun (2025) [8] found that the pilot zone policy improved firm labor allocation efficiency, with AI technology penetration being a key transmission pathway. Zhao and Liu (2025) [18] further found a U-shaped relationship between AI penetration rate and firm innovation efficiency: the efficiency-enhancing effect only fully manifests after the penetration rate crosses a certain threshold, and this process is positively moderated by regional human capital levels. From a broader industrial chain perspective, Wang et al. (2026) [10] found that the AI application of core firms generates bilateral spillover effects in the supply chain, causing the markup distribution to shift from polarization toward coordination, with a particularly large enhancing effect on upstream firms. Taken together, the micro-level impacts of AI policy do exist, but their direction and magnitude vary across firms; which firms can benefit and to what extent is largely constrained by firms' own technological foundations and regional conditions.

The second strand involves the mechanisms through which AI and digital technologies affect firm markup. Markup is a commonly used indicator for measuring firm pricing power and market power (De Loecker et al., 2020) [25], and can also be used to observe whether involutionary competition exists among firms. Existing research shows that the application of AI and digital technologies improves firm markup along multiple dimensions. Yin et al. (2026a) [12] pointed out that human–AI collaboration can reduce redundant investment, eliminate inefficient capacity, advance core technology breakthroughs, and improve supply–demand matching, thereby promoting markup increases through multiple pathways. Yin et al. (2026b) [13] found that computing power deployment helps firms overcome the middle-technology trap and absorb excess capacity, exerting a restraining effect on price-cutting involution. Li et al. (2023) [5], Yu et al. (2025) [15], and Zou (2024) [19], from the respective perspectives of firm digital transformation, digital talent allocation, and digital economy development, have all observed similar positive effects. Among them, Zou (2024) [19] also found that the digital economy, by raising markups, indirectly promotes firms' ascent along the global value chain.

In addition, some studies have found nonlinear relationships between the two. Deng et al. (2024) [4] observed a U-shaped trajectory between digital transformation and markup, with low-level digitalization actually depressing markup, which only turns upward after crossing the inflection point. Dai et al. (2023) [3], in contrast, reported an inverted U-shaped relationship. These two seemingly contradictory conclusions actually point to the same fundamental fact: the effect of AI technology on markup does not occur spontaneously, nor is it a simple linear increase; rather, it is constrained by stage-specific conditions and particular contexts. The finding of Bai and Yu (2021) [9] further indicates that digital economy development may depress firm price markups, because intensified competition and impeded cost pass-through jointly outweigh efficiency improvements. This result does not genuinely contradict the aforementioned positive effects; rather, it precisely reveals the dual nature of the digital economy's impact, whose net effect depends on which force—efficiency enhancement or competition intensification—proves more powerful, and this balance of forces is highly correlated with firms' own conditions, market structures, and institutional environments.

For the above contradictory phenomena, absorptive capacity theory provides an effective explanatory framework. Cohen and Levinthal (1990) [24] pointed out that a firm's ability to utilize external knowledge presupposes that it possesses a relevant stock of prior knowledge. Building on this, Zahra and George (2002) [27] further distinguished between the two concepts of potential absorptive capacity and realized absorptive capacity. According to this theory, even if industrial policies are open to all firms, the distribution of policy dividends cannot be equal. Only those firms with sufficient absorptive capacity can transform the opportunities brought by external technological shocks into tangible market gains. This assessment has been validated in empirical research in the AI domain. Zhao and Liu (2025) [18] found that the effect of AI penetration rate on innovation efficiency is positively moderated by human capital levels, and human capital is precisely a key component of absorptive capacity. Yin et al. (2024) [14] provided more direct threshold effect evidence: only after AI penetration crosses a certain scale and achieves a transition from labor substitution to human–machine collaboration do productivity-enhancing effects manifest. These

findings indicate that the heterogeneity of AI technology effects is not attributable to random factors but is endogenous to systematic differences in firm absorptive capacity.

However, the measurement of absorptive capacity is itself a challenge. Early research used R&D intensity as a proxy variable (Cohen and Levinthal, 1990) [24], but this indicator cannot capture firms' attention to specific technologies and the direction of knowledge accumulation. In recent years, methods based on annual report text analysis have provided a new solution path. By extracting the frequency of terms related to specific technologies in corporate annual reports, one can measure firms' strategic attention to and cognitive depth in a particular technology domain (Zhao and Liu, 2025) [18]. Complementary to text-based indicators are patent-based measurement methods. Patent data reflect firms' actual innovation output and knowledge stock in specific domains (Breschi et al., 2003; Leten et al., 2007) [23][26]. The research of Xu et al. (2019) [11] and Zhang et al. (2018) [17] indicates that in rapidly iterating fields such as AI technology, the focus of technology penetration may be more critical than the breadth of general technology diversification. The finding of Chen and Ou (2026) [2] from a data asset perspective also supports this assessment: the productivity effect of data factors is more pronounced in non-state-owned and non-high-tech industries, indicating that firms with weaker traditional resource endowments may also achieve latecomer advantages through focused attention on specific technology domains.

Synthesizing the above literature, it can be concluded that existing research has not yet placed AI policy shocks, firm AI technology penetration, and markup within a unified causal identification framework. Research on the micro-level heterogeneity of AI policy effects provides identification strategies and a quasi-natural experimental basis for this paper, but the existing literature has mostly focused on labor allocation efficiency and innovation efficiency, with relatively little direct investigation of firms' pricing power in product markets. Research on the determinants of markup provides rich mechanism clues for this paper, but most studies focus on the spontaneous diffusion of technology rather than the structural shocks of specific policies. Absorptive capacity theory provides theoretical support for understanding effect heterogeneity, but how to measure absorptive capacity in the AI domain *ex ante* and incorporate it into a causal identification framework remains an unresolved problem. This paper advances on the basis of the above literature: taking the AI Pilot Zone policy as a quasi-natural experiment, measuring firm AI technology penetration through AI-related term frequency in annual reports, employing the DDD method to identify the heterogeneity of AI policy effects across firms with different AI technology penetration levels, and testing the innovation output channel and the moderating role of R&D absorptive capacity.

2.2 Research Hypotheses

Based on the above review of the literature, this paper constructs a theoretical analysis framework from the perspective of absorptive capacity. The AI Pilot Zone policy creates an external technological knowledge pool for firms through channels such as reducing the cost of AI technology acquisition, cultivating industrial ecosystems, and attracting talent agglomeration. However, according to the core insight of absorptive capacity theory, a firm's ability to derive benefits from external knowledge is constrained by its prior accumulation of relevant knowledge. This determines that the impact of AI policy on firm performance will not be homogeneous, but will exhibit systematic differences conditioned on firms' technological readiness.

With respect to the average effect of AI policy on firm markup, the AI Pilot Zone policy may exert influence through two channels. The first is the technology empowerment channel. AI Pilot Zones, by providing high-quality data resources and technology services (Yun et al., 2026) [16], reduce firms' AI technology application costs and enhance their product differentiation capability and production efficiency (Yu et al., 2025; Yin et al., 2026a) [15][12], thereby strengthening firms' pricing power in product markets. The second is the market structure channel. The application of AI technology may intensify market competition and reinforce winner-take-all effects (Autor et al., 2020) [22], exerting pressure on the pricing space of some firms (Bai and Yu, 2021) [9]. The net effect of the two channels depends on which force dominates. In the Chinese context, the AI Pilot Zone policy is still in the early stages of diffusion, and the technology

empowerment effect may be stronger than the competitive squeeze effect. Accordingly, this paper proposes:

Hypothesis 1: The AI Pilot Zone policy has a positive effect on firm markup.

The journey from AI technology absorption to the realization of commercial value places relatively high demands on firms' own technological foundations and degree of attention. AI technology penetration reflects the level of AI technology attention and application presented in firms' annual reports, and also serves as a basis for assessing firms' potential absorptive capacity. After the policy is implemented, high-penetration firms more rapidly initiate AI technology innovation, constructing market barriers through differentiated products and strengthened technological advantages, thereby obtaining greater markup enhancement potential (Wang et al., 2026) [10]. In contrast, firms with lower penetration, lacking corresponding knowledge preparedness, face more obstacles in absorbing and utilizing policy resources. Theoretically, this moderating effect can be explained through the innovation output channel: high-penetration firms increase their AI-related patent applications and build competitive barriers through sustained innovation accumulation. Accordingly, this paper proposes:

Hypothesis 2: Firm AI technology penetration positively moderates the impact of AI policy shocks on firm markup, and this moderating effect is reflected at the innovation output end.

The impact of AI policy on firm operational efficiency may exhibit more pronounced capability threshold characteristics. Existing research has shown that the release of AI technology dividends has threshold effects (Yin et al., 2024) [14]; only when firms possess a minimum level of absorptive capacity can external technological shocks be tangibly translated into efficiency improvements. Firm R&D intensity, as a measure of general technological absorptive capacity, may play a moderating role in the process through which AI policy affects operational efficiency. Even when the average effect of AI policy on asset turnover is not significant, high-R&D firms may still benefit from the policy. Accordingly, this paper proposes:

Hypothesis 3: Firm R&D intensity positively moderates the impact of AI policy shocks on firm operational efficiency. This effect may exist independently of the average effect of AI policy on asset turnover.

3. Research Design

3.1 Model Specification

Baseline Model (DID). This paper exploits the exogenous variation in the location and timing of the "New-Generation Artificial Intelligence Development Plan" pilot zones to construct the following difference-in-differences model:

$$Markup_{ict} = \beta_0 + \beta_1(Post_{ct} \times Treat_c) + \gamma X_{ict} + \mu_i + \lambda_t + \epsilon_{ict}$$

where i , c , and t denote firm, city, and year, respectively; $Markup_{ict}$ is the natural logarithm of firm markup; $Post_{ct}$ is a dummy variable for the post-policy period (taking the value of 1 for the year in which the pilot zone is approved and all subsequent years); $Treat_c$ is a treatment group dummy variable (taking the value of 1 if the firm's registered location is in an AI Pilot Zone city); X_{ict} is a vector of firm-level control variables; μ_i and λ_t are firm fixed effects and year fixed effects, respectively. Standard errors are clustered at the firm level.

Moderation Effect Model (DDD). To identify the moderating effect of AI technology penetration, the triple interaction term is introduced:

$$Markup_{ict} = \phi_0 + \phi_1(Post_{ct} \times Treat_c) + \phi_2(Post_{ct} \times Treat_c \times AI_wf_std_{it}) + \phi_3 AI_wf_std_{it} + \gamma X_{ict} + \mu_i + \lambda_t + \epsilon_{ict}$$

where $AI_wf_std_{it}$ is the standardized measure of firm AI technology penetration, and ϕ_2 is the core identification parameter: if $\phi_2 > 0$, then firms with higher AI technology penetration obtain a larger markup increase from the AI policy.

To test the supplementary moderating effect of R&D intensity, the dependent variable is replaced with asset turnover (ATO_{it}), and the moderating variable is replaced with R&D intensity (RD_{it} , standardized):

$$ATO_{ict} = \psi_0 + \psi_1(Post_{ct} \times Treat_c) + \psi_2(Post_{ct} \times Treat_c \times RD_{it}) + \psi_3RD_{it} + \gamma X_{ict} + \mu_i + \lambda_t + \epsilon_{ict}$$

Event Study Model. To test the parallel trends assumption and characterize the dynamic evolution of the policy effect, an event study model incorporating relative-time dummy variables is constructed:

$$Markup_{ict} = \alpha + \sum_{k=-5, k \neq -1}^5 \theta_k D_{it}^k + \gamma X_{ict} + \mu_i + \epsilon_{ict}$$

where D_{it}^k is an indicator variable for the relative time to the policy (with $k = -1$ as the reference period). This model includes firm fixed effects to control for time-invariant firm characteristics.

Mechanism Testing Model. To test the transmission channels of the moderating effect of AI technology penetration, the impact of the triple interaction term is estimated with firm AI invention patent applications ($\ln_AI_invention_{it}$), R&D intensity (RD_{it}), and cash flow ratio ($Cash_flow_{it}$) as dependent variables, respectively:

$$M_{ict} = \eta_0 + \eta_1(Post_{ct} \times Treat_c) + \eta_2(Post_{ct} \times Treat_c \times AI_wf_std_{it}) + \eta_3AI_wf_std_{it} + \gamma X_{ict} + \mu_i + \lambda_t + \epsilon_{ict}$$

where M_{ict} takes the above mechanism variables in turn.

3.2 Variable Definitions

Dependent Variables. This paper involves two dependent variables. The first is firm markup ($\ln_markup$), obtained through ACF two-stage estimation of a translog production function, with its natural logarithm taken in the regression. The second is firm operational efficiency (ATO), measured by asset turnover, i.e., the ratio of operating revenue to total assets. The above data are all sourced from the pre-computed datasets of the CSMAR database.

Core Explanatory Variable. The core explanatory variable is $Post \times Treat$, a binary dummy variable for the AI Pilot Zone policy. $Treat$ indicates whether the firm's registered location is in an AI Pilot Zone city (1 if yes, 0 otherwise); $Post$ indicates the policy implementation timing (1 for the year of approval and all subsequent years, 0 for years prior to approval). The product of the two variables yields $Post \times Treat$: 0 before approval, 1 after approval for firms in Pilot Zone cities, and always 0 for non-Pilot Zone cities. According to MOST announcements, the AI Pilot Zones were established in three batches: Beijing and Shanghai were approved in 2019; 12 cities including Tianjin, Shenzhen, Hangzhou, Hefei, Chengdu, Xi'an, Chongqing, Guangzhou, Wuhan, Jinan, Suzhou, and Changsha were approved in 2020; and Shenyang, Harbin, and Zhengzhou were approved in 2023.

Moderating Variables. There are two moderating variables. The first is AI technology penetration (AI_word_freq), obtained through text analysis of listed firms' annual reports. Specifically, the frequency of terms such as "artificial intelligence," "machine learning," "deep learning," "neural network," "natural language processing," "computer vision," "intelligent robot," "large language model," and "generative AI" as a proportion of total word count is computed, thereby measuring firms' degree of attention to and level of application of AI technologies. This variable is standardized before entering the regression and is denoted as AI_wf_std . The second is R&D intensity ($RD_intensity$), measured as the ratio of firm R&D expenditure to total assets, standardized and denoted as RD_std .

Mechanism Variables. AI invention patent applications ($\ln_AI_invention$): the natural logarithm of the annual number of firm AI-related invention patent applications plus one, sourced from the patent database of the China National Intellectual Property Administration. This variable serves as the dependent variable in the first step of the mechanism test, used to examine whether the moderating effect of AI technology penetration operates through increased AI patent applications. In addition, R&D intensity ($RD_intensity$) and cash flow ratio ($Cashflow$) also serve as dependent variables in the mechanism tests, respectively investigating the innovation input and cash flow channels; their measurement methods have been described in the moderating variable and control variable sections above.

Control Variables. At the firm level: firm size (Size), measured as the natural logarithm of total assets; debt-to-asset ratio (Lev), measured as the ratio of total liabilities to total assets; return on assets (ROA), measured as the ratio of net profit to total assets; firm age ($\ln_FirmAge$), measured as the natural logarithm of the number of years since establishment; ownership type (SOE), taking the value of 1 for state-owned enterprises and 0 otherwise; Tobin's Q (TobinQ), measured as the ratio of firm market value to asset replacement cost; cash flow ratio (Cashflow), measured as the ratio of net cash flow from operating activities to total assets. At the city level (controlled in robustness checks): city economic development level ($\ln_CityGDP$), measured as the natural logarithm of city GDP; industrial structure (IndustryStructure), measured as the share of secondary sector value added in regional GDP; education level ($\ln_Education$), measured as the natural logarithm of the number of regular institutions of higher education; Internet penetration rate (Internet), measured as the ratio of Internet broadband access subscribers to the registered household population. In addition, dummy variables for contemporaneous S&T policies are controlled in robustness checks: the Broadband China pilot policy (DID_Broadband), the Big Data Comprehensive Pilot Zone policy (DID_BigData), and the Public Data Open Platform policy (DID_OpenData); each policy takes the value of 1 for the year of city approval and all subsequent years.

3.3 Data Sources and Descriptive Statistics

The sample of this paper consists of panel data of A-share listed manufacturing firms from 2014 to 2024. Firm financial and corporate governance data are sourced from the CSMAR database; markup data are from the pre-computed ACF estimation dataset; AI patent data are from the patent database of the China National Intellectual Property Administration; annual report text data are from the CNRDS database; and AI Pilot Zone DID data are compiled from MOST announcements. The sample is screened according to the following criteria: (1) financial firms are excluded; (2) ST and PT firms are excluded; (3) only manufacturing firms (industry code beginning with C) are retained; (4) continuous variables are winsorized at the 1st and 99th percentiles. The final sample comprises 3,943 firms with 27,696 firm-year observations, of which 13,141 are non-missing observations for the dependent variable ($\ln_markup$). Table 1 presents the descriptive statistics of the main variables. In addition, this paper conducts a balance test on key variables between the treatment and control groups during the pre-policy period (2014–2018); see Appendix Table 1 for details.

Table 1: Descriptive Statistics of Main Variables

Variable	Observations	Mean	Std. Dev.	Min	Median	Max
$\ln_markup$	13,141	0.135	0.102	0.004	0.112	0.531
ATO	27,691	0.627	0.344	0.099	0.561	2.096
Post_Treat	27,696	0.231	0.422	0.000	0.000	1.000
AI_word_freq	27,653	0.362	0.507	0.000	0.191	3.215
$\ln_AI_invention$	27,696	1.182	1.505	0.000	0.693	9.027
$RD_intensity$	27,073	4.757	4.912	0.000	3.480	25.821

Variable	Observations	Mean	Std. Dev.	Min	Median	Max
Size	27,696	22.032	1.180	19.156	21.889	26.085
Lev	27,696	0.384	0.192	0.059	0.372	0.887
ROA	27,696	0.041	0.067	-0.208	0.042	0.210
In_FirmAge	27,696	1.382	0.076	1.080	1.397	1.550
SOE	25,455	0.234	0.423	0.000	0.000	1.000
TobinQ	27,351	2.085	1.236	0.864	1.682	7.687
Cashflow	27,696	0.050	0.066	-0.141	0.049	0.226

4. Empirical Results

4.1 Baseline Regression Results

Table 2 reports the stepwise regression results of the baseline DID and DDD models. Column 1, the parsimonious DID specification including only firm and year fixed effects, shows that the $Post \times Treat$ coefficient is negative and fails the significance test. In Column 2, after introducing firm-level control variables, the $Post \times Treat$ coefficient becomes 0.0055 and reaches the 5% significance level, thereby supporting Hypothesis 1. In economic terms, the AI Pilot Zone policy increases firm markup by approximately 0.55% on average, equivalent to 4.07% of the dependent variable mean. After adding the triple interaction term—the DDD model specification—in Column 3, the coefficient of $Post \times Treat \times AI_wf_std$ is 0.0038, significant at the 5% level, providing support for Hypothesis 2's expectation of a moderating effect. This means that for each one-standard-deviation increase in AI technology penetration, the markup-enhancing effect of the AI policy increases by an additional 0.38 percentage points, approximately 2.81% of the dependent variable mean. Column 4 further controls for city-level variables, after which the triple interaction coefficient declines to 0.0035 with a p-value rising to 0.071, reaching marginal significance at the 10% level. This change indicates, to some extent, that city-level characteristics partially absorb the AI policy effect—that is, there exists a certain degree of correlation between the economic and educational resource endowments of Pilot Zone cities and firms' AI technology penetration.

Table 2: Baseline DID and DDD Regression Results

Variable	(1) DID Without Controls	(2) DID + Firm Controls	(3) DDD + Firm Controls	(4) DDD + All Controls
Post_Treat	-0.0008 (0.0037)	0.0055** (0.0028)	0.0052* (0.0027)	0.0022 (0.0029)
AI_wf_std			-0.0004 (0.0013)	-0.0007 (0.0015)
Post×Treat×AI_wf_std			0.0038** (0.0019)	0.0035* (0.0019)
Firm Controls	No	Yes	Yes	Yes
City Controls	No	No	No	Yes
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	13,141	12,118	12,105	10,401

Variable	(1) DID Without Controls	(2) DID + Firm Controls	(3) DDD + Firm Controls	(4) DDD + All Controls
Number of Firms	2,749	2,669	2,669	2,314
R ²	0.000	0.443	0.443	0.436

Notes: ***, **, * denote significance at the 1%, 5%, and 10% levels, respectively; values in parentheses are firm-level clustered robust standard errors. Firm control variables include Size, Lev, ROA, In_FirmAge, SOE, TobinQ, and Cashflow. City control variables include In_CityGDP, In_Education, IndustryStructure, and Internet. Subsequent tables use the same control variables unless otherwise noted.

4.2 Event Study and Dynamic Effects

To test the parallel trends assumption and characterize the dynamic evolution of the policy effect, this paper employs the event study method, taking one year before policy implementation ($t = -1$) as the reference period and estimating the relative effects for each period before and after the policy. The event study uses a firm fixed effects model without year fixed effects to avoid perfect collinearity with the event dummy variables; the control group consists of firms never treated by the policy.

Table 3: Event Study Results (D₋₁ as Reference Period)

Relative Period	Coefficient	Std. Err.	p-value
D _{m5} (t-5)	-0.0236	0.0045	0.000
D _{m4} (t-4)	-0.0215	0.0043	0.000
D _{m3} (t-3)	-0.0299	0.0037	0.000
D _{m2} (t-2)	-0.0072	0.0033	0.027
D _{p0} (t=0)	0.0001	0.0027	0.967
D _{p1} (t+1)	-0.0064	0.0031	0.038
D _{p2} (t+2)	-0.0123	0.0033	0.000
D _{p3} (t+3)	-0.0147	0.0036	0.000
D _{p4} (t+4)	-0.0223	0.0041	0.000
D _{p5} (t+5)	-0.0199	0.0053	0.000

Notes: N = 12,118, R² = 0.444. Includes firm fixed effects and firm-level control variables. The control group consists of firms never treated by the policy.

The results in Table 3 show that the coefficients for all pre-policy periods (D_{m5} through D_{m2}) are negative and statistically significant, indicating that the treatment group and the never-treated group already exhibited differential markup trends prior to policy implementation. This result implies that the conventional DID parallel trends assumption is difficult to satisfy in this sample.

The identification strategy of this paper is based on the triple-difference method. Relative to conventional DID, DDD does not require parallel trends in the outcome variable between the treatment and control groups; rather, it requires that firms at different levels of AI technology penetration exhibit parallel trends in markup changes before and after the policy. To test this assumption, this paper splits the sample into high-penetration and low-penetration groups by the median of AI technology penetration, separately estimates the event study model for each group, and examines whether the two groups exhibit systematic differences in pre-policy coefficients. Figure 1 reports the grouped event study results.




Figure 1: DDD Parallel Trends Test

Figure 1 shows that for the high AI technology penetration group, the coefficients for all pre-policy periods fail the significance test, indicating that this group exhibited no significant markup trend differences prior to policy implementation. For the low AI technology penetration group, marginally significant negative coefficients exist in the earlier pre-policy periods ($t = -5$ to $t = -3$), but their direction is consistent with the high-penetration group, and the post-policy coefficients for both groups are not significant and close to zero. The above results indicate that the parallel trends assumption on which the DDD identification strategy relies is broadly satisfied; in particular, for the high-penetration firm group, which contributes substantially to the DDD interaction term, no significant differences in pre-policy trends are present.

4.3 Placebo Test

To rule out the possibility that unobservable random factors generate spuriously significant results, this paper conducts a placebo test by randomly assigning treatment timing; the results are presented in Figure 2. The specific procedure is as follows: while maintaining the binary structure of treatment and control groups unchanged, a pseudo policy timing is randomly assigned to each firm, with the pseudo timing randomly drawn from the five actual policy years from 2019 to 2023. Based on the random timing, a pseudo triple interaction term is reconstructed, and the DDD model is repeatedly estimated 2,000 times to obtain the empirical distribution of the pseudo coefficients.

After 2,000 random simulations, the mean of the pseudo triple interaction coefficients is approximately 0.0024, with the distribution clearly concentrated near zero. The actual estimate of 0.0024 lies at the 49.4th percentile of this distribution, corresponding to an empirical p-value of 0.506. The true effect does not significantly exceed the pseudo effects under randomization, indicating that the DDD coefficient is not driven by random noise in the timing of treatment assignment. In conjunction with the parallel trends evidence from Figure 1 and the extensive robustness checks in Section 4.4, the identification strategy yields credible estimates of the moderating effect of AI technology penetration.




Figure 2: Placebo Test

4.4 Robustness Checks

To ensure the reliability of the research conclusions, this paper conducts eight robustness checks across four dimensions—variable measurement, policy interference exclusion, sample composition testing, and model specification adjustment. Table 4 summarizes all test results. In addition, after adding city-level control variables in Column 4 of the baseline regression, the triple interaction coefficient is 0.0035 ($p = 0.071$), consistent in direction with the baseline results, indicating that the conclusions are insensitive to city-level characteristics.

First, substitution of core variable measurement methods. To test whether the conclusions depend on particular variable construction methods, this paper conducts verification along two directions. First, R&D intensity is used in place of AI technology penetration as the moderating variable. R&D intensity and AI technology penetration capture firms' technological absorptive capacity from different angles: the former reflects firms' general level of R&D investment, while the latter reflects firms' degree of attention to the specific domain of AI. The regression results show that the triple interaction coefficient is 0.0068 with a p-value of 0.060, maintaining a positive direction with a larger effect magnitude. Second, the original markup level value is used in place of log-transformed markup as the dependent variable. The results show that the triple interaction coefficient is 0.0046 with a p-value of 0.051, consistent in direction. The above two tests indicate that the conclusions of this paper are not driven by particular variable construction methods.

Second, exclusion of interference from contemporaneous S&T policies. Multiple AI-related policies existed during the sample period and could potentially interfere with the identification of the independent effect of the Pilot Zone policy. To address this issue, this paper additionally includes in the baseline regression the DID variables for three contemporaneous S&T policies—Broadband China, Big Data Comprehensive Pilot Zones, and Public Data Open Platforms—and re-estimates the model. The results show that the triple interaction coefficient is 0.0039 with a p-value of 0.041, broadly consistent with the baseline regression, indicating that the moderating effect of the Pilot Zone policy is not attributable to the influence of other policies.

Third, testing the sensitivity of sample composition. The estimation results of the baseline regression may be influenced by specific sample points, and this paper tests this along four dimensions. The first dimension is the exclusion of the four directly administered municipalities—Beijing, Shanghai, Tianjin, and Chongqing. These cities are mostly among the first-batch or key pilot zones, with markedly higher concentrations of S&T resources than other regions, and may exert a substantial pull on the overall estimation results. After exclusion, the triple interaction coefficient is 0.0037 with the p-value rising to 0.132, indicating that the municipality samples contribute substantially to statistical significance. The second dimension is the exclusion of the year 2020 to remove the disturbance of the COVID-19 pandemic on firm operations and policy transmission. The results show that the triple interaction coefficient rises to 0.0043 with a p-value of 0.037, becoming even more significant, suggesting that the pandemic factor introduced some interference in the estimation of the policy effect, and that the policy effect becomes clearer after its exclusion. The third dimension is the application of 2.5% winsorization to the dependent variable and the moderating variable to control for the influence of extreme values; the results show that the triple interaction coefficient is 0.0041 with a p-value of 0.039, leaving the conclusions unchanged. The fourth dimension is the retention of only the subsample of patent-holding firms; the triple interaction coefficient is then 0.0032 with a p-value of 0.096, maintaining the same direction but with reduced significance.

Fourth, alternative fixed effects structures. The baseline model specification employs firm fixed effects and year fixed effects. To test the sensitivity of the conclusions to the form of fixed effects, this paper replaces year fixed effects with industry-year interactive fixed effects, allowing different industries to have their own time trends, while using firm dummy variables to control for firm fixed effects. Under this more stringent specification, the triple interaction coefficient is 0.0038 with a p-value of 0.007, with significance rising to the 1% level, indicating that the baseline conclusions are insensitive to the specific form of fixed effects.

Table 4: Robustness Checks

Variable	(1) Alternative Moderator	(2) Alternative DV	(3) Control Other Policies	(4) Excl. Municipalities	(5) Excl. COVID Year	(6) 2.5% Winsorization	(7) Patent Subsample	(8) Industry×Year FE
Post×Treat×Moderator	0.0068* (0.0036)	0.0046* (0.0024)	0.0039** (0.0019)	0.0037 (0.0025)	0.0043** (0.0021)	0.0041** (0.0020)	0.0032* (0.0019)	0.0038*** (0.0014)
Observations	12,005	12,105	12,105	10,631	10,567	12,105	11,282	12,105
R ²	0.452	0.426	0.443	0.453	0.440	0.440	0.439	0.868
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
City Controls	No	No	No	No	No	No	No	No
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
Industry×Year FE	No	No	No	No	No	No	No	Yes

Notes: (1) R&D intensity replaces AI_wf_std as the moderating variable. (2) Original markup level replaces log-transformed markup as the dependent variable. (3) Additionally controls for three DID variables: Broadband China, Big Data Pilot Zones, and Public Data Open Platforms. (4) Excludes the four directly administered municipalities. (5) Excludes the 2020 COVID-19 sample. (6) 2.5% winsorization applied to the dependent and moderating variables. (7) Retains only the patent-holding firm subsample. (8) Industry×year fixed effects replace year fixed effects; firm dummies replace firm fixed effects (OLS estimation).

4.5 Heterogeneity Analysis

This paper conducts heterogeneity analysis along three dimensions—ownership type, industry technology intensity, and region—to examine the differential performance of the moderating effect of AI technology penetration across different contexts. Table 5 summarizes the grouped regression results.

First, heterogeneity by firm ownership type. This paper splits the sample into state-owned enterprises and private enterprises for grouped regressions. With respect to the significance of the moderating effect, the triple interaction coefficient for private enterprises is 0.0052, significant at the 5% level with a p-value of 0.016. For state-owned enterprises, the coefficient is negative and not significant, with a p-value of 0.474. In terms of coefficient magnitude, the private enterprise coefficient is approximately 0.0076 higher than that of state-owned enterprises. A possible explanation for this difference is that private enterprises are more agile in market competition and can more rapidly translate AI technology penetration into actual product differentiation and pricing power. State-owned enterprises, constrained by organizational inertia and non-market decision-making mechanisms, find it difficult to efficiently convert AI technology attention into market advantages even if they possess it.

Second, heterogeneity by industry technology intensity. This paper splits the sample into technology-intensive industries and traditional industries. The regression coefficient for technology-intensive industries is 0.0034, and that for traditional industries is 0.0025; both coefficients are in the same positive direction. In terms of significance, neither group reaches the 10% significance level, and the between-group difference also fails the statistical test. This indicates that the moderating effect of AI technology penetration is broadly present across industries with different technology levels, but the intensity differences have not yet reached statistical significance.

Third, heterogeneity by region. This paper splits the sample into eastern and central-western regions based on firm location for grouped regressions. The triple interaction coefficient for the eastern region is 0.0063, significant at the 5% level with a p-value of 0.047. The coefficient for the central-western region is 0.0035, failing the significance test with a p-value of 0.149. In terms of coefficient magnitude, the effect size in the eastern region is approximately 1.8 times that in the central-western region. The advantages of the eastern region in AI talent agglomeration, digital infrastructure, and the degree of marketization provide firms with more favorable external conditions. However, it should be noted that the eastern region subsample size is only 1,335, far smaller than the 10,770 in the central-western region, and the statistical inference for this result should be interpreted with caution.

Table 5: Heterogeneity Analysis

Variable	(1) SOE	(2) Private	(3) Tech-Intensive	(4) Traditional	(5) Eastern	(6) Central-Western
Post×Treat×AI_wf_std	-0.0024	0.0052**	0.0034	0.0025	0.0063**	0.0035
	(0.0033)	(0.0022)	(0.0022)	(0.0031)	(0.0031)	(0.0024)
Observations	2,736	9,369	8,217	3,888	1,335	10,770
R ²	0.457	0.436	0.437	0.477	0.382	0.450
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes
City Controls	No	No	No	No	No	No
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Technology-intensive industries are defined as manufacturing industry codes C26 through C40. In the ownership-type grouping, the SOE variable is excluded from the control variables to avoid collinearity.

4.6 Mechanism Tests

To verify the transmission channels of the moderating effect of AI technology penetration, this paper conducts tests along three dimensions: innovation output, R&D investment, and cash flow. All three channels adopt a unified identification strategy, with each mechanism variable serving as the dependent variable and the triple interaction term serving as the core explanatory variable. Table 6 presents the corresponding results. First, the innovation output channel. The regression results with firm AI invention patent applications as the dependent variable show that the coefficient of $Post \times Treat \times AI_wf_std$ is 0.0477, significant at the 5% level (p -value = 0.026), indicating that for each one-standard-deviation increase in AI technology penetration, the policy shock drives an increase in AI patent applications of approximately 4.8%. Hypothesis 2's expectation regarding the innovation output channel receives empirical support. Second, the R&D investment channel. The regression results with firm R&D intensity as the dependent variable show that the triple interaction coefficient is 0.1728, reaching marginal significance at the 10% level (p -value = 0.061), indicating that the AI policy shock prompts firms with higher AI technology penetration to further increase their R&D investment. Third, the cash flow channel. The regression results with firm operating cash flow ratio as the dependent variable show that the triple interaction coefficient is positive but fails the significance test (p -value = 0.135), indicating that the moderating effect of the AI policy does not operate through the channel of cash flow improvement; firms are more inclined to reinvest policy resources in technological innovation activities rather than convert them into short-term cash flow.

Table 6: Mechanism Tests

Variable	(1) AI Patent Applications	(2) R&D Intensity	(3) Cash Flow Ratio
Post_Treat	0.1061***	0.2981**	0.0054***
	(0.018)	(0.121)	(0.001)
AI_wf_std	0.0184	0.0333	-0.0001
	(0.014)	(0.096)	(0.001)
Post×Treat×AI_wf_std	0.0477**	0.1728*	0.0018
	(0.026)	(0.061)	(0.135)
Observations	25,096	24,547	25,096
Firm FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Firm Controls	Yes	Yes	Yes

Notes: The dependent variables for the three regressions are, respectively, log of AI invention patent applications plus one, R&D intensity, and cash flow ratio. Control variables are the same as in the baseline model.

4.7 Further Tests of the Moderation Effect

In addition to the moderating effect of AI technology penetration on markup, this paper further examines the moderating role of R&D intensity in the operational efficiency effect of the AI policy, thereby testing Hypothesis 3. Table 7 reports the corresponding regression results. The DID estimate in Model 1 shows that the average effect of the AI policy on asset turnover is negative and fails the significance test (p -value = 0.347), indicating that, at the average level, the AI Pilot Zone policy did not directly drive improvements in firm operational efficiency. The DDD estimation results in Model 2 show that the coefficient of $Post \times Treat \times RD_std$ is 0.0195, significant at the 1% level (p -value = 0.0002), thus verifying Hypothesis 3.

This means that for each one-standard-deviation increase in R&D intensity, the asset-turnover-enhancing effect of the AI policy increases by approximately 1.95 percentage points. The above results jointly indicate that the operational efficiency effect of the AI policy is highly dependent on firms' own R&D investment levels: for firms with insufficient R&D investment, the AI policy can hardly be directly translated into improvements in operational efficiency; for high-R&D firms, the AI policy, by strengthening their technology absorption and transformation capabilities, significantly promotes efficiency gains. This result is consistent with the theoretical expectations of absorptive capacity theory.

Table 7: R&D Intensity Moderates the Operational Efficiency Effect

Variable	Model 1 DID	Model 2 DDD
Post_Treat	-0.0077	-0.0076
	(0.0081)	(0.0083)
RD_std		-0.0698***
		(0.0059)
Post×Treat×RD		0.0195***
		(0.0053)
Firm Controls	Yes	Yes
Firm FE	Yes	Yes
Year FE	Yes	Yes
Observations	25,111	24,562
Number of Firms	3,849	3,765
R ²	0.216	0.248

Notes: The dependent variable is asset turnover (ATO), measuring firm operational efficiency. Control variables are the same as in the baseline model.

5. Conclusions and Policy Recommendations

5.1 Main Conclusions

Leveraging the National New-Generation Artificial Intelligence Innovation Development Pilot Zones as a quasi-natural experiment, using panel data of listed manufacturing firms from 2014 to 2024, and employing DID and DDD methods, this paper incorporates firm AI technology penetration into the policy effect evaluation framework and tests the applicability and boundaries of absorptive capacity theory in explaining the micro-level effects of S&T industrial policies. The main findings are as follows.

First, the National New-Generation AI Innovation Development Pilot Zone policy has, on average, led to an increase in firm markup, but the distribution of this increase is uneven. From the involution perspective, whether the AI Pilot Zone policy can alleviate the dilemma of firms improving efficiency without raising prices depends on firms' own technological readiness. Firms with higher AI technology penetration can more quickly convert policy dividends into competitive advantages in the market, while firms with weak technological reserves benefit notably less. This divergence is also reflected at the ownership type and regional levels: private enterprises and firms in eastern regions exhibit more pronounced performance, indicating that absorptive capacity is not merely an internal technological condition of firms but is also related to the institutional environment and regional endowments. Overall, the National New-Generation

AI Innovation Development Pilot Zone policy has, in an average sense, improved the situation of firms enhancing efficiency without increasing prices, but exhibits pronounced differentiation at the micro level: firms with ample reserves have used it to consolidate their market positions, while firms with insufficient reserves face pressure instead.

Second, the average impact of the AI policy on firm operational efficiency is not significant, but firms with higher R&D intensity have indeed experienced efficiency improvements, indicating that absorptive capacity plays a screening role in policy transmission. Firms with insufficient absorptive capacity, even when facing the same policy shock, may find it very difficult to obtain tangible benefits—policy dividends may bypass them. With respect to mechanisms, firms with higher AI technology penetration are indeed more active in innovation investment and patent applications, but there exists a time lag between innovation output and pricing power. For technological innovation to truly affect market competitiveness, it must pass through the stages of productization, commercialization, and scaling; if problems arise at any stage, the policy effect may stall at the level of patent counts without reaching the level of market power.

Third, AI technology penetration and R&D intensity respectively act on two different output dimensions—pricing power and operational efficiency—and these two capabilities cannot be simply substituted for one another. Firms need to accumulate the corresponding technological capabilities according to their own target directions. The above results support the explanatory power of absorptive capacity theory for S&T industrial policies: simply promoting technology diffusion without simultaneously building absorptive capacity may cause policies to objectively amplify rather than narrow gaps among firms.

5.2 Policy Recommendations

Based on the above research findings, this paper proposes the following policy recommendations.

First, establish a tiered support system based on firm technological capabilities. AI policy dividends are predominantly captured by firms with higher technological reserves, while firms with weak foundations find it difficult to benefit. It is recommended that a firm technological capability assessment component be embedded within the AI Pilot Zone policy framework, and differentiated support schemes be implemented based on firms' AI technology penetration and R&D investment levels. For small and medium-sized enterprises with weak technological foundations, upfront technical training, talent sharing, and pilot project subsidies can be provided to help them cross the minimum capability threshold for absorbing AI technologies. For technologically leading firms, emphasis can be placed on industry–university–research platform construction and generic technology R&D support to amplify their technology spillover effects. At the same time, it is recommended that AI industrial policy and R&D subsidy policy form a temporal sequence: first enhance firms' absorptive capacity through basic R&D investment, then release the dividends of technology diffusion through AI policy.

Second, strengthen the role of market mechanisms in AI technology diffusion. Private enterprises demonstrate stronger capability in converting policy dividends into market advantages; it is recommended that policy design make greater use of market-based instruments such as government procurement, computing power vouchers, and technical standards guidance to reduce the transaction costs and entry barriers for firms applying AI technologies. At the same time, administrative tilts toward specific ownership types should be reduced, so that technological capability rather than ownership type becomes the primary basis for firms to obtain policy support.

Third, construct inter-regional technological collaboration and transfer mechanisms. The regional differences in AI policy effects indicate a tendency for technology dividends to concentrate in the east. It is recommended that the construction of AI Pilot Zones in central and western regions be accompanied by the introduction of technological collaboration resources from eastern regions, transferring the technological experience and managerial capabilities of pioneer regions to central and western firms through channels such as industrial alliances, technology export, and talent exchange. This will help mitigate inter-regional capability gaps during the early stages of technology diffusion and prevent the gap

in digital economy development from further widening due to the asymmetric distribution of policy dividends.

5.3 Research Limitations

This paper has the following limitations, which also point toward directions for future research. First, there remains room for improvement in the measurement precision of absorptive capacity. This paper employs AI-related term frequency in annual reports as a proxy for AI technology penetration; this indicator reflects firms' strategic attention to AI technology rather than the actual depth of their technology application. As annual report texts are externally disclosed documents, firms may overestimate or underestimate their actual level of investment in AI, and the mapping relationship between the word-frequency indicator and real technological capability is subject to information noise. Future research could integrate multi-dimensional information such as AI job recruitment data, AI software procurement records, and AI-related investments to construct a more comprehensive absorptive capacity measurement system. Second, there is room for improvement in both causal identification and mechanism testing. At the identification strategy level, the event study results show that trend differences existed between the treatment group and the never-treated group prior to the policy. Although the identification assumptions of the DDD method do not require the treatment and control groups to satisfy level parallel trends, whether firms at different levels of AI technology penetration exhibit systematic differences in pre-policy markup trends still requires more rigorous testing to confirm. At the mechanism level, this paper verifies the promotional effect of the policy shock on AI patent output, but the conversion pathway from patents to market pricing power has not yet received statistical verification. On the one hand, this may be because the observation window after the policy is relatively short, and the conversion from patent applications to competitive advantage in product markets requires time to materialize; on the other hand, it may also be because patents exert influence through multiple indirect channels that are difficult to fully identify through the two-stage regression approach. Subsequent research could consider employing methods such as synthetic control, interaction-weighted estimation, or distributed lag models to conduct more detailed analysis of the above issues.

Notes:

- ① "Involution" in the economic context refers to the phenomenon whereby firms improve efficiency but are unable to translate efficiency gains into profit margins or pricing power; its essence is a decline in firm markup (see De Loecker et al., 2020).
- ② "New-Generation Artificial Intelligence Development Plan" (Guofa [2017] No. 35), State Council, July 2017.
- ③ See the MOST official replies on supporting Beijing and Shanghai in establishing National New-Generation AI Innovation Development Pilot Zones (Guoke Han Gui [2019] No. 27 and No. 28), as well as subsequent batches of Pilot Zone approval documents.
- ④ For the estimation method of markup, see De Loecker, J., Eeckhout, J., & Unger, G. (2020). The rise of market power and the macroeconomic implications. *Quarterly Journal of Economics*, 135(2), 561–644.
- ⑤ For the classic exposition of absorptive capacity, see Cohen, W. M., & Levinthal, D. A. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly*, 35(1), 128–152.

References

- [1] Cai Y Z, Chen N. Artificial intelligence and high-quality growth and high-quality employment under the new technological revolution [J]. *Journal of Quantitative & Technological Economics*, 2019, 36(5): 3–22. (In Chinese)
- [2] Chen Z S, Ou G G. How can new-type production factors resolve the "productivity paradox"? Micro evidence from firm data assets [J]. *Modern Economic Research*, 2026, (5): 1–13. (In Chinese)

- [3] Dai X, Ma H W, Zhang E Z. Can digital transformation necessarily increase firm markup? [J]. Journal of Financial Research, 2023, (5): 134–151. (In Chinese)
- [4] Deng S J, Zhang Q, Chen X X. The path of firm transformation in the digital era: A markup perspective on China's listed industrial firms [J]. Journal of Finance and Economics, 2024, 50(6): 153–168. (In Chinese)
- [5] Li H, Liang B Q, Cao Z X. Research on the impact of digital transformation on firm cost markup [J]. Journal of Hunan University (Social Sciences), 2023, 37(4): 65–76. (In Chinese)
- [6] Yang C, Gao W L, Chang S W. How does generative AI technology affect firm total factor productivity? Innovation measurement based on the technology patent classification system [J/OL]. Journal of Quantitative & Technological Economics, 1–27 [2026-06-15]. (In Chinese)
- [7] Liu W L, Cheng Q. Digital industry penetration, global production networks, and asymmetric technology spillovers [J]. China Industrial Economics, 2023, (3): 96–114. (In Chinese)
- [8] Mao J H, Sun J J. How does "AI+" affect firm labor allocation efficiency? Empirical evidence from National New-Generation AI Innovation Development Pilot Zones [J/OL]. Northwest Population Journal, 1–16 [2026-06-15]. (In Chinese)
- [9] Bai P W, Yu L. Digital economy development and firm price markup: Theoretical mechanisms and empirical facts [J]. China Industrial Economics, 2021, (11): 59–77. (In Chinese)
- [10] Wang X D, Liu D, Wang S Y. Research on the impact of AI supply chain spillovers on firm markup distribution [J/OL]. Chinese Journal of Management, 1–10 [2026-06-15]. (In Chinese)
- [11] Xu L Y, Zeng D M, Li J. The impact of knowledge network centrality and knowledge diversification on firm ambidextrous innovation performance [J]. Management Review, 2019(7): 110–121. (In Chinese)
- [12] Yin Q, Yu X, Fan L B. Research on the impact of human–AI collaboration on firm cost markup: With a discussion on breaking "involutionary" competition [J]. Inquiry into Economic Issues, 2026, (5): 120–138. (In Chinese)
- [13] Yin Q, Fan L B, Niu B. Alleviating "involutionary" price-cutting competition: The impact of computing power deployment on firm cost markup [J/OL]. Foreign Economics & Management, 1–18 [2026-06-15]. (In Chinese)
- [14] Yin Y H, Miao Y, Sun X D. Intelligent technology penetration, labor market transformation, and the release of technology dividends [J]. Journal of Southwest Minzu University (Humanities and Social Science), 2024, 45(7): 102–115. (In Chinese)
- [15] Yu D F, Zhang Y, Song Q. Digital human capital investment and firm price markup: Evidence from recruitment big data [J]. China Industrial Economics, 2025, (7): 62–80. (In Chinese)
- [16] Yun Q Y, Chen Y C, Ge J F. Data factors promote the development of intelligent information technology industries: From the perspective of digital technology penetration [J]. Shanghai Journal of Economics, 2026, (3): 116–128. (In Chinese)
- [17] Zhang Q L, Shi J J, Liu C L. Technology diversification, industry competitive interaction, and ambidextrous innovation capability [J]. Foreign Economics & Management, 2018(9): 71–83. (In Chinese)
- [18] Zhao J B, Liu Z. Research on the impact of AI penetration rate on firm innovation efficiency [J]. Economic Perspectives, 2025, (5): 55–69. (In Chinese)
- [19] Zou J. Digital economy, firm markup, and global value chain position ascent [J]. Statistics & Decision, 2024, 40(4): 140–145. (In Chinese)
- [20] Acemoglu D, Restrepo P. The race between man and machine: Implications of technology for growth, factor shares, and employment [J]. American Economic Review, 2018, 108(6): 1488–1542.

- [21] Aghion P, Jones B F, Jones C I. Artificial intelligence and economic growth [R]. National Bureau of Economic Research, 2017.
- [22] Autor D, Dorn D, Katz L F, et al. The fall of the labor share and the rise of superstar firms [J]. The Quarterly Journal of Economics, 2020, 135(2): 645–709.
- [23] Breschi S, Lissoni F, Malerba F. Knowledge-relatedness in firm technological diversification [J]. Research Policy, 2003, 32(1): 69–87.
- [24] Cohen W M, Levinthal D A. Absorptive capacity: A new perspective on learning and innovation [J]. Administrative Science Quarterly, 1990, 35(1): 128–152.
- [25] De Loecker J, Eeckhout J, Unger G. The rise of market power and the macroeconomic implications [J]. The Quarterly Journal of Economics, 2020, 135(2): 561–644.
- [26] Leten B, Belderbos R, Van Looy B. Technological diversification, coherence, and performance of firms [J]. Journal of Product Innovation Management, 2007, 24(6): 567–579.
- [27] Zahra S A, George G. Absorptive capacity: A review, reconceptualization, and extension [J]. Academy of Management Review, 2002, 27(2): 185–203.

Appendix

Appendix Table 1 reports the balance test results for key variables between the treatment and control groups during the pre-policy period (2014–2018). The results show that 7 of the 10 variables exhibit significant differences at the 5% level: the treatment group is significantly higher than the control group in markup (ln_markup), AI word frequency (AI_word_freq), R&D intensity (RD_intensity), and Tobin's Q (TobinQ), and significantly lower in firm size (Size), firm age (ln_FirmAge), and cash flow (Cashflow). These between-group differences are the endogenous outcome of the siting logic of the AI Innovation Development Pilot Zones—namely, that Pilot Zones preferentially select cities with rich S&T resources and strong innovation capacity, leading firms located in these cities to differ naturally in technological attention and operational characteristics. However, this imbalance does not constitute a fundamental threat to the DDD identification strategy: first, all regressions in this paper include firm fixed effects, such that time-invariant firm characteristics (including level differences between the treatment and control groups) are fully absorbed; second, DDD identification relies on differences among firms with different AI penetration levels within the same treatment group, rather than on level comparisons between the treatment and control groups.

Appendix Table 1: Pre-Policy Balance Test (2014–2018)

Variable	Treatment Group Mean	Control Group Mean	Difference	p-value
ln_Markup	0.1442	0.1302	0.0141***	0.000
AI_word_freq	0.4964	0.2893	0.2071***	0.000
RD_intensity	6.0670	4.0739	1.9932***	0.000
Size	22.1309	22.1677	-0.0368**	0.045
Lev	0.4304	0.4298	0.0006	0.828
ROA	0.0433	0.0449	-0.0017*	0.075
ln_FirmAge	1.3478	1.3578	-0.0100***	0.000
SOE	0.2692	0.2545	0.0147	0.119

Variable	Treatment Group Mean	Control Group Mean	Difference	p-value
TobinQ	2.3788	2.1847	0.1941***	0.000
Cashflow	0.0419	0.0482	-0.0062***	0.000

Notes: The treatment group consists of firms located in AI Pilot Zone cities; the control group consists of firms in other cities. p-values are from two-sample t-tests (allowing for heteroskedasticity). ***, **, * denote significance at the 1%, 5%, and 10% levels, respectively.