Impact of AI Adoption Announcements on Stock Returns in the Chinese Market: An Event Study Perspective

by

Yijing Chen

An honors thesis submitted in partial fulfillment

of the requirements for the degree of

Bachelor of Science

Business and Economics Honors Program

NYU Shanghai

May 2026

Professor Marti G. Subrahmanyam Professor Christina Wang Professor Wendy Jin	Professor Christina Wang
Faculty Advisers	Thesis Adviser

Abstract

This paper investigates the stock market reaction to announcements of artificial intelligence (AI) adoption by Chinese-listed firms. Using an event study approach combined with natural language processing (NLP) classification, the study examines abnormal returns across different event windows and analyzes detailed AI application scenarios. The results reveal that AI-related announcements generate mixed short-term reactions but exhibit a positively skewed distribution in the medium term. Firms with stronger AI commercialization potential and smaller market capitalization experience more significant positive effects, while high-volatility firms are penalized over longer horizons. Internal AI development is associated with short-term gains but does not sustain in the long run. These findings provide new evidence on how emerging technologies are priced in financial markets and highlight the importance of distinguishing between different types of AI strategies.

Keywords: Artificial Intelligence, Event Study, Stock Market Reaction, Abnormal Returns, Natural Language Processing, Chinese A-share Market

Content

1. Introduction

With the rapid advancement of conversational artificial intelligence (AI) models such as ChatGPT and DeepSeek, AI-related themes have increasingly influenced stock market dynamics, particularly within the technology sector. Recently, AI has evolved from a broad technological narrative into a practical corporate strategy, with firms adopting AI tools to improve products, customer service, operational efficiency, and strategic decision-making. In the Chinese market, this trend has been accompanied by a growing number of listed companies publicly disclosing their adoption or development of AI-related technologies. In this paper, AI adoption announcements refer to firm-level disclosures indicating the integration, application, or development of AI technologies in corporate activities, including internal operations, external products and services, or industry-specific AI solutions.

Such announcements may affect stock prices by influencing investor expectations about future profitability, growth potential, and technological competitiveness. Compared with conventional financial disclosures, AI-related announcements are often forward-looking, strategically framed, and difficult for investors to value precisely. As a result, they may trigger heterogeneous market reactions through changes in investor sentiment, capital flows, and perceived firm value. This raises an important empirical question: how do AI adoption announcements affect stock prices, and do these effects vary across time horizons and firm characteristics? More specifically, this study asks whether AI-related disclosures generate abnormal returns in the short run, whether these reactions persist or gradually reverse over time, and whether the magnitude of the market response depends on factors such as firm size, volatility, commercialization potential, and the textual content of the announcement itself.

A large body of literature examines how financial markets respond to corporate announcements using event study methodologies. The event study framework provides a standard approach to evaluate the impact of new information on firm value by analyzing abnormal stock returns around announcement dates (MacKinlay, 1997; Kothari and Warner, 2007). Prior research has documented that stock prices react to a wide range of firm-specific events, including earnings releases, mergers and acquisitions, and corporate investment decisions, suggesting that financial markets incorporate new information into prices in a timely manner. In particular, a growing strand of research focuses on market reactions to technological innovation and corporate investment in new technologies. Existing studies have examined the valuation effects of R&D investments, patent filings, and IT-related announcements, showing that markets tend to respond positively to signals of future growth potential (Chan et al., 2001; Hirshleifer et al., 2013). For instance, IT investment announcements have been found to significantly affect firm valuation, reflecting investor expectations about technological adoption and productivity improvements (Takeda et al., 2021).

More recently, financial research has increasingly incorporated textual analysis and natural language processing (NLP) techniques to extract structured information from unstructured financial data. Prior studies demonstrate that textual information, such as media tone and corporate disclosures, can influence investor sentiment and predict stock returns (Tetlock, 2007; Tetlock et al., 2008). Additionally, domain-specific tools, such as the financial sentiment dictionary developed by Loughran and McDonald (2011), have enhanced the accuracy of textual analysis in financial contexts.

Despite the extensive literature on corporate announcements and technological innovation, there remains limited empirical evidence on how AI adoption announcements, as distinct from general innovation or technology news, affect stock returns. This gap is particularly important because AI-related disclosures are highly heterogeneous: some emphasize internal technological capability, while others focus on external adoption or revenue-generating applications, which may lead investors to respond differently.

To address these questions, this study adopts a three-step empirical strategy. First, firm-level AI adoption announcements are systematically processed using a natural language processing (NLP)-based classification framework, which extracts structured information on the source, application scope, and implementation stage of AI-related disclosures. This step allows the construction of both the main analytical sample and a set of text-derived explanatory variables. Second, an event study approach is employed to examine market reactions around announcement dates by estimating abnormal returns (AR) and cumulative abnormal returns (CAR) over multiple event windows. This analysis provides initial evidence on whether AI adoption announcements generate significant stock price responses and how these responses evolve over time. Third, cross-sectional regression models are used to investigate the determinants of abnormal returns, linking CAR to NLP-derived variables and firm characteristics. Additional specifications incorporate interaction terms and aggregated variables to explore heterogeneity in market responses and to identify the underlying mechanisms through which different types of AI adoption are valued by investors.

Therefore, this study can contribute to the literature in several ways. First, it provides new evidence on the impact of AI adoption announcements on the Chinese A-share market, a setting where AI-related narratives and retail investor participation are especially prominent. Second, by applying natural language processing (NLP) techniques to classify AI-related announcements into distinct categories, the paper offers a more nuanced analysis of how investors price different types of AI strategies. Third, the study combines event-study analysis with cross-sectional regressions to examine not only whether abnormal returns exist but also which firm characteristics and announcement features help explain their magnitude and persistence. Overall, the paper aims to shed light on how emerging AI strategies are interpreted and valued in equity markets.

2. Data and Preprocess

2.1 Data Source

The primary data source for this study is the Wind Terminal, from which firm-level AI-related announcements and stock market data were collected. The sample period for AI adoption announcements covers February 2025 to August 2025, focusing on Chinese A-share listed firms. This time window was selected because AI adoption announcements were relatively sparse before February 2025, partly due to the still-limited practical capabilities and commercial readiness of large language model technologies. Beginning in February 2025, however, a sharp increase in relevant corporate disclosures emerged, reflecting a broader acceleration in firm-level AI adoption and market attention. To ensure comprehensive coverage and reduce potential sample selection bias, the data collection process incorporated announcements related to a wide range of major domestic large language model (LLM) platforms, including Doubao, Kimi, DeepSeek, Tencent Hunyuan, Huawei Pangu, and others. This approach helps mitigate the risk of relying on a single model, platform, or keyword systematically biasing the sample toward specific types of firms or announcements. Moreover, the sample was truncated in August 2025 to ensure sufficient post-event trading data for the event study analysis, particularly for long-term event windows.

2.2 Sample Construction and Data Cleaning

For each event, the dataset obtained from the Wind Terminal includes firm-level announcement texts, event dates, and corresponding stock price information required for the event study analysis. To accurately identify the event date, the timing of each AI adoption announcement was cross-verified using multiple sources, including the publication timestamp of the news article, the effective or disclosed date mentioned in the announcement text, the release date on the firm's official website, and the posting date on the firm's official WeChat account. When discrepancies existed across sources, the earliest verifiable disclosure date was retained as the event date. This approach was adopted to better capture the point at which the information first became available to the market and to reduce potential bias arising from delayed reporting or information leakage. In addition, firm-level financial characteristics and control variables were merged into the sample to support the subsequent cross-sectional regression analysis. Because not all AI-related announcements necessarily reflect substantive firm-level AI adoption, the announcement texts were further processed using an NLP-based classification framework. Details of the classification process are provided in Section 2.3. To reduce potential confounding effects from concurrent earnings disclosures and other financially material corporate information, events for which the event date occurred within seven days of the firm's most recent financial reporting date were excluded from the final sample.

2.3 NLP-Based Classification

To systematically distinguish different types of AI-related disclosures, this study applies a large language model (LLM)-based classification framework to the announcement texts. The classification serves two purposes. First, it is used to determine whether a disclosure reflects substantive firm-level AI adoption and should therefore be included in the main analytical sample. Second, it is used to generate structured text-based variables that capture the strategic content of each announcement for the subsequent cross-sectional analysis.

Specifically, the announcement texts were first classified using a locally deployed Qwen2.5:7b model, guided by a rule-based prompt designed to extract economically meaningful event characteristics. A central objective of this procedure was to distinguish substantive AI adoption disclosures from broad, forward-looking, or purely conceptual references to AI, as well as announcements focused solely on computing infrastructure or supportive technical capabilities without specifying concrete firm-level actions or application scenarios. The model converted unstructured announcement text into a set of structured labels describing the nature, source, scope, functional role, and implementation stage of the disclosed AI activity. These labels were then used both to determine whether an announcement belonged to the main analytical sample and to construct text-derived explanatory variables for the subsequent empirical analysis. After applying the sample selection criteria, including the identification of substantive AI adoption announcements and the requirement of sufficient pre-event trading data, the final sample consists of 104 firm-level events. Table 1 summarizes the definitions and analytical roles of the NLP-derived classification labels.

To improve classification reliability, the same announcement texts were also independently classified using a second LLM-based system, ChatGPT 5.1, which served as an external reference for comparison. The final labels were then determined through manual review and reconciliation, especially in cases where the two model outputs differed or where the announcement language was ambiguous. This hybrid classification pipeline was designed to improve consistency, reduce misclassification, and ensure that the final event sample reflected economically meaningful AI adoption behavior rather than broad or purely conceptual AI narratives. The full prompt and detailed classification rules are provided in Appendix A.

Table 1

Definition of NLP-Derived Variables

Variable / Label	**Possible Values**	**Definition**
is_ai_related	True / False	Indicates whether the announcement describes a concrete AI application scenario rather than a purely conceptual or strategic reference to AI.
ai_source	external / internal / mixed / none	Identifies whether the disclosed AI technology is externally adopted, internally developed, or both.
application_scope	internal / external / mixed / none	Captures whether the AI application is used for internal firm operations, external products / services, or both.
internal_support	operational_efficiency / data_intelligence	Functional category of internal AI use, test only when *application_scope* is internal or mixed
external_support	product_ai_enhancement / customer_service_ai / industry_solution	Functional category of external AI use, tested only when *application_scope* is external or mixed
adoption_stage	completed / in_development / future_plan / none	Whether the AI adoption is completed, under development, or only planned
main_sample	yes / no	Whether the announcement reflects substantive, ongoing, or completed AI adoption and is included in the final sample

2.4 Variable Construction

This subsection describes the construction of the key variables used in the empirical analysis. The dependent variable in this study is cumulative abnormal return (CAR), measured over multiple event windows, including short-term post-event windows such as [0, 5] and medium-term windows such as [0, 20]. To reduce the influence of extreme observations, cumulative abnormal return (CAR) measures are winsorized at the 1st and 99th percentiles. The detailed calculation of abnormal returns is presented in Section 3.

The main explanatory variables are derived from the NLP-based classification described in Section 2.3. These variables capture different dimensions of firm-level AI adoption, including the source of AI technology, the scope of application, and the stage of adoption. In particular, several derived explanatory variables are constructed from the classification labels summarized in Table 1. The variable *internal_ai_model* is coded as one when the announcement indicates either internal AI development or a mixed strategy combining internal and external AI adoption, thereby capturing whether the firm possesses internal AI capability. The variable *external_revenue_ai* is constructed as a composite indicator of AI-related commercialization and equals one when the announcement reflects either a product AI enhancement or an industry solution in external-facing applications. This variable is intended to capture AI use cases that are more directly associated with revenue generation or external value creation.

A set of firm-level control variables is also included in the analysis. These variables include profitability (ROA), valuation (price-to-book ratio), leverage, liquidity, and stock return volatility. To mitigate potential endogeneity concerns, all financial variables are measured using the most recently available quarterly information prior to the event date. Volatility is calculated as the annualized standard deviation of daily stock returns over the [-60, -5] pre-event estimation window. Firm size is controlled for using market capitalization. Given the highly right-skewed distribution of market capitalization in the event sample, firms are categorized into size terciles based on the 33rd and 67th percentiles of market capitalization within the sample. This approach avoids imposing arbitrary absolute size thresholds and ensures relatively balanced group sizes for statistical inference.

In addition, selected interaction terms are constructed for supplementary analyses, such as the interaction between *external_revenue_ai* and firm profitability (ROA), in order to examine whether the market response to AI-related commercialization varies with underlying firm characteristics.

3. Methodology

3.1 Event Study Design

This study employs a standard event study methodology to examine stock market reactions to AI adoption announcements, following the framework developed by MacKinlay (1997). The event study framework evaluates whether abnormal returns occur around the announcement date, reflecting the market's response to new information.

The abnormal return for firm *i* on day *t* is defined as:

$$AR_{i,t} = R_{i,t} - \hat{R}_{i,t}$$

where $R_{i,t}$ is the actual stock return and $\hat{R}_{i,t}$ is the expected return estimated using asset pricing models. In the baseline specification, expected returns are estimated using the Fama–French three-factor model, which accounts for market, size, and value factors. For robustness, expected returns are also estimated using the Capital Asset Pricing Model (CAPM). Both models are estimated over an estimation window of [-200, -5] trading days prior to the event date.

Cumulative abnormal returns (CAR) are then calculated as:

$$CAR_{i,[t1,t2]} = \sum AR_{i,t}$$

The analysis focuses on multiple event windows, including short-term windows such as [0, 5] and medium-term windows such as [0, 20], in order to capture both immediate and delayed market reactions.

3.2 Cross-Sectional Regression

To examine the determinants of cross-sectional variation in market reactions, this study conducts regression analysis using cumulative abnormal returns (CAR) as the dependent variable. The baseline regression specification is as follows:

$$CAR_i = \alpha + \beta X_i + \gamma Controls_i + \epsilon_i$$

where CAR_i represents the cumulative abnormal return for firm *i* over a given event window. The vector X_i includes key explanatory variables derived from the NLP-based classification, such as indicators for internal AI capability (*internal_ai_model*), AI-related commercialization (*external_revenue_ai*), and the stage of AI adoption.

The regression also includes a set of firm-level control variables, including profitability (ROA), valuation (price-to-book ratio), leverage, liquidity, stock return volatility, and firm size. These variables are included to account for heterogeneity in firm characteristics that may influence market reactions. In supplementary specifications, time fixed effects are also introduced as a robustness check to account for common shocks associated with announcement timing.

In addition, interaction terms are introduced in selected specifications to examine whether the market response to AI-related announcements varies with firm characteristics. For example, the interaction between *external_revenue_ai* and ROA is used to test whether the valuation effect of AI commercialization depends on firm profitability.

4 Results

4.1 Benchmark Comparison and Event-Time Patterns

This subsection examines the market reaction to AI adoption announcements under different benchmark models. Specifically, cumulative abnormal returns (CAR) are estimated using both the CAPM and the Fama-French benchmarks, allowing an assessment of whether the observed results are sensitive to the choice of expected return model.

Table 2

Summary Statistics of Cumulative Abnormal Returns by Benchmark

Benchmark	Variable	Mean	Median	Std. Dev.	25th pct	75th pct	t-stat
CAPM	CAR [0,5]	0.056	0.008	0.149	-0.025	0.089	3.853***
	CAR [0,20]	0.041	0.018	0.173	-0.047	0.094	2.435**
Fama-French (3-factor)	CAR [0,5]	0.044	0.005	0.136	-0.025	0.068	3.329***
	CAR [0,20]	0.023	0.008	0.164	-0.059	0.078	1.452

Note: This table reports summary statistics of cumulative abnormal returns (CAR) across different event windows under the CAPM and Fama–French three-factor benchmarks. The t-statistic tests whether the mean CAR is significantly different from zero. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 2 reports summary statistics of abnormal returns across different event windows under the two benchmark specifications. As shown in the table, the mean CAR over the [0, 5] window is positive and statistically significant under both models, indicating that AI adoption announcements are associated with economically meaningful short-term stock price reactions. However, the evidence becomes weaker over longer horizons. While the mean CAR over the [0, 20] window remains positive, it is no longer statistically significant under the Fama–French benchmark, suggesting that the initial market reaction may not fully persist over time. In addition, the difference between mean and median CAR indicates a right-skewed distribution, implying that the average effect is partly driven by a subset of firms experiencing relatively large positive returns.

Figure 1 further illustrates the event-time dynamics of cumulative abnormal returns under the two benchmark models.

Figure 1

Event-Time CAR Trajectories Under CAPM and Fama–French Benchmarks

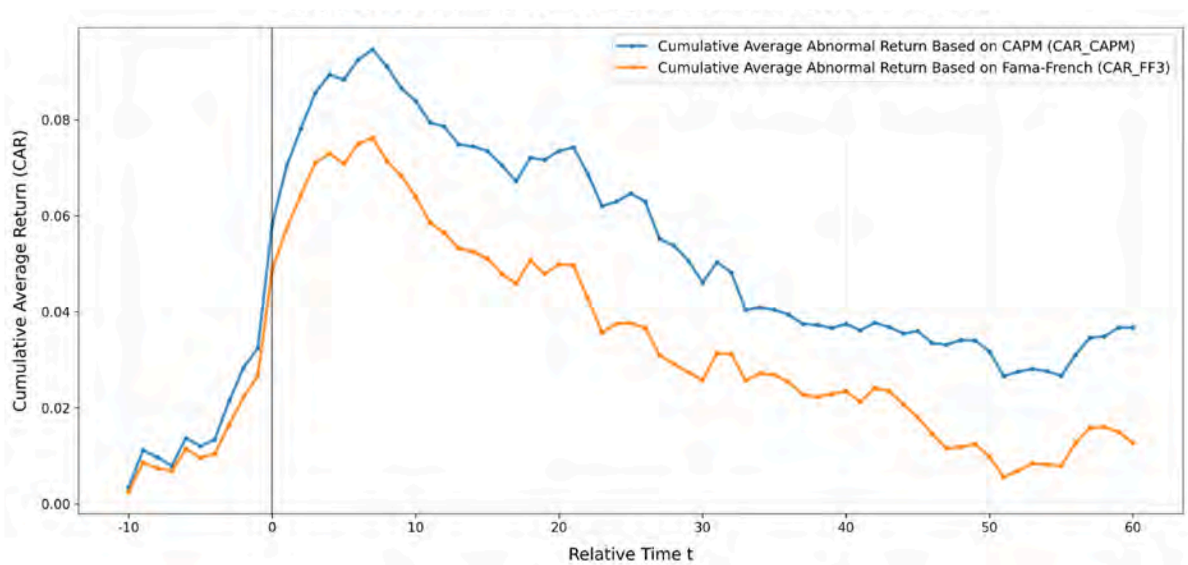

Overall, both specifications reveal meaningful event-time variation, but the Fama–French benchmark appears to capture the short-term abnormal return dynamics more clearly, particularly within the [0, 5] post-event window. This suggests that controlling for additional risk factors such as size and value may improve the measurement of expected returns in the context of Chinese A-share firms, where cross-sectional heterogeneity is substantial.

In addition, both benchmark models show a noticeable increase in cumulative abnormal returns during the [-5, 0] pre-event window. This pattern may indicate that part of the information was incorporated into prices before the formal disclosure date, potentially reflecting information leakage, early market anticipation, or gradual diffusion of announcement-related signals prior to the official release. Given these patterns, the Fama-French benchmark is used as the baseline specification in the subsequent analysis, while the CAPM-based results are retained as a robustness check.

4.2 Distribution of Cumulative Abnormal Returns

After establishing the benchmark model for expected return estimation, this subsection examines the distribution of cumulative abnormal returns (CAR) across firm-level AI adoption announcements. Understanding the cross-sectional distribution of CAR is useful for assessing whether market reactions are broadly concentrated, symmetric, or characterized by substantial heterogeneity across firms.

Figures 2 and 3 present the distributions of CAR under the Fama–French benchmark for the [0, 5] and [0, 20] event windows, respectively. Several patterns emerge:

First, the short-term distribution of CAR[0, 5] is centered close to zero, with the largest concentration of observations falling in the range of [-5%, 0%], followed by [0%, 5%]. This suggests that the immediate market response to AI adoption announcements is generally modest rather than uniformly strong, and that many announcements do not generate large abnormal returns in the very short run. At the same time, the distribution is not symmetric: while most firms experience relatively small short-term reactions, a subset of firms exhibits noticeably larger positive abnormal returns, producing a visible right tail.

Second, the distribution of CAR[0, 20] shifts moderately toward positive values. Compared with the short-term window, the mass of observations becomes more concentrated in the [0%, 5%] range, while the right tail remains pronounced. This pattern suggests that the market response to AI adoption announcements may unfold more gradually over time, with some firms receiving a more sustained positive reassessment in the medium term. In other words, the effect of AI-related disclosures does not appear to be fully incorporated immediately upon announcement for all firms.

Figure 2

Distribution of CAR Over the [0, 5] Event Window

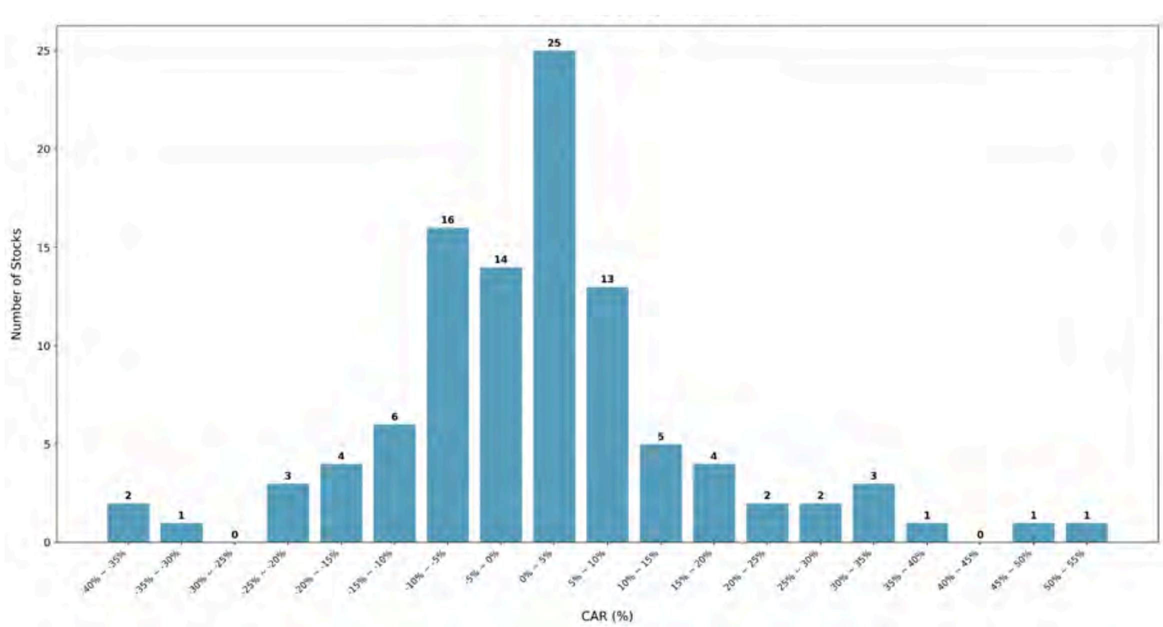

Note: This figure shows the distribution of cumulative abnormal returns (CAR) over the [0, 5] trading-day event window under the Fama–French benchmark. CAR values are winsorized at the 1st and 99th percentiles to reduce the influence of extreme observations.

Figure 3

Distribution of CAR Over the [0, 20] Event Window

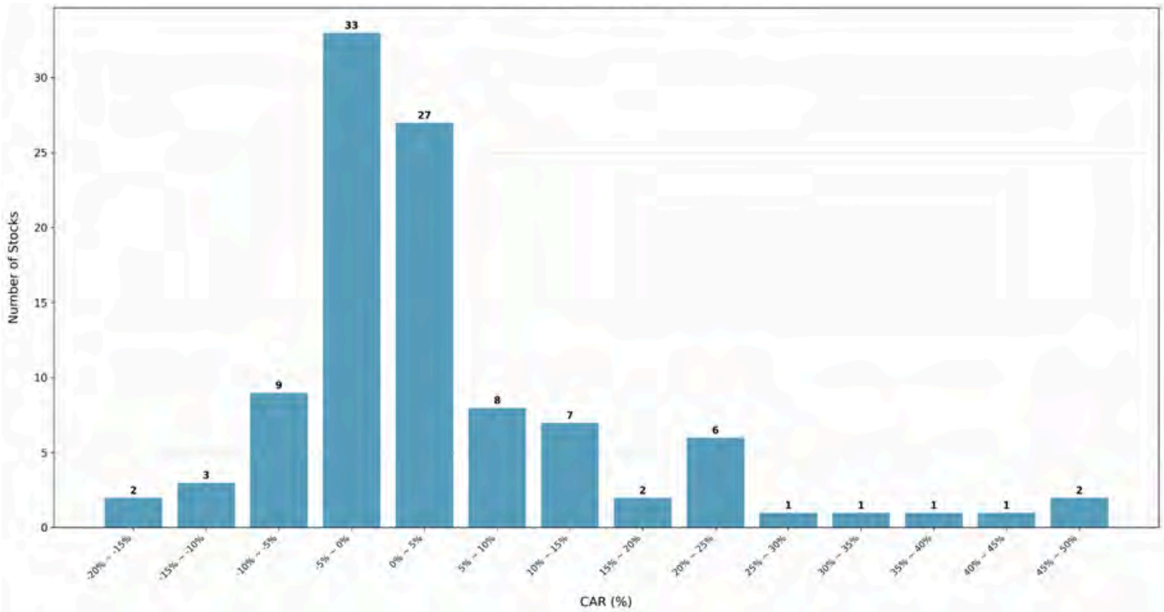

Note: This figure shows the distribution of cumulative abnormal returns (CAR) over the [0, 20] trading-day event window under the Fama–French benchmark. CAR values are winsorized at the 1st and 99th percentiles to reduce the influence of extreme observations.

Third, both distributions exhibit substantial cross-sectional heterogeneity, particularly in the medium-term window. While a large share of firms cluster around small abnormal returns, a smaller group of firms experiences substantially larger positive or negative reactions. This heterogeneity is consistent with the idea that investors do not respond uniformly to all AI adoption announcements. Instead, the market reaction likely depends on differences in announcement content, firm characteristics, and the perceived economic value of the disclosed AI strategy. To reduce the influence of extreme observations while preserving the full sample, CAR measures are winsorized at the 1st and 99th percentiles, as described in Section 2.4.

4.3 Cross-Sectional Regression Results

This subsection examines whether the cross-sectional variation in cumulative abnormal returns can be explained by differences in announcement content and firm characteristics. In particular, the regression analysis tests whether the market reacts differently to AI adoption announcements depending on the firm's internal AI capability, the commercial relevance of the announced application, the implementation stage, and financial characteristics. The detailed regression results for the short-term and medium-term event windows are reported in Tables 3 and 4, respectively.

Overall, the regression results suggest that the market does not respond uniformly to AI adoption announcements. Instead, investors appear to differentiate between disclosures that signal commercial potential, implementation credibility, and firm-level risk exposure.

In the short-term window [0, 5], several announcement-related variables are positively associated with abnormal returns, as shown in Table 3. In particular, the variable *ext_application* is significantly positive in the short-term specifications, indicating that the market reacts more positively when firms disclose AI adoption in ways that are directly connected to external-facing use cases, rather than to broad, purely conceptual, or exclusively internal AI narratives. Consistent with this interpretation, the variable *external_revenue_ai*, a composite indicator capturing AI applications linked to product enhancement or industry solution, is also significantly positive, with an estimated effect in the range of approximately 7% to 9%.

This finding suggests that the market places particular value on AI disclosures that are more directly associated with external value creation or potential revenue generation. To further unpack the content behind these external-facing AI applications, additional decomposition analyses were conducted, with detailed results reported in Appendix Tables A1 and A2. These results show that the positive effect of *external_revenue_ai* is driven primarily by announcements classified as industry_solution, while *product_ai* and *service_ai* do not exhibit comparable statistical significance. This suggests that investors respond more strongly to AI disclosures framed as industry-specific commercial solutions, rather than to more general product enhancement or service-oriented AI applications.

Another notable short-term result is the interaction term between *external_revenue_ai* and firm profitability (*ROA*), which is significantly negative in the [0, 5] window, as reported in Table 3. This finding suggests that the short-term market premium associated with AI-related commercialization is weaker among more profitable firms. One possible interpretation is that AI-related revenue narratives generate stronger abnormal reactions when they are disclosed by firms for which such announcements represent a more meaningful growth or revaluation signal, rather than by firms that are already financially strong and more fully priced by the market.

Moreover, a subsample analysis restricted to firms that had already adopted external AI technologies shows that *internal_ai_model*, an indicator equal to one when the announcement reflects either internal AI development or a mixed strategy combining internal and external AI adoption, remains positively associated with short-term abnormal returns. This pattern suggests that, even in an environment where external large language model adoption has become widespread, the market may still place additional value on firms that also demonstrate internal AI development capability. Detailed results are reported in Appendix Tables A1 and A2.

Firm size also plays an important role. The coefficient on *SmallCap* is significantly negative in both the [0, 5] and [0, 20] windows across several specifications in Tables 3 and 4, with an estimated effect ranging from approximately -10% to -7% relative to the benchmark middle-size group.

Table 3

Cross-Sectional Regression Results for CAR[0, 5]

Variables	(1)	(2)	(3)	(4)
External_Revenue_AI	—	—	0.077** (0.029)	0.087*** (0.029)
External_Revenue_AI × ROA	—	—	—	-0.017** (0.007)
Ext_Application	0.078*** (0.028)	0.078*** (0.028)	—	—
Comp_Stage	0.059 (0.038)	0.070 (0.044)	0.058 (0.040)	0.064 (0.039)
Comp_Stage × SmallCap	—	-0.051 (0.090)	—	—
SmallCap	-0.073** (0.032)	-0.028 (0.086)	-0.074** (0.033)	-0.093*** (0.033)
LargeCap	-0.068* (0.040)	-0.073* (0.039)	-0.067* (0.040)	-0.073* (0.039)
ROA	-0.003 (0.004)	-0.003 (0.004)	-0.004 (0.004)	0.008 (0.006)
Leverage	0.000 (0.001)	0.000 (0.001)	0.000 (0.001)	-0.000 (0.001)
PB	-0.004 (0.004)	-0.003 (0.004)	-0.002 (0.004)	-0.001 (0.004)
Liquidity	-0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)
Volatility	0.018 (0.074)	0.014 (0.072)	0.006 (0.075)	0.005 (0.073)
Constant	-0.030 (0.079)	-0.025 (0.063)	0.009 (0.064)	0.012 (0.062)

Note: This table reports cross-sectional regression results using cumulative abnormal returns over the [0, 5] event window as the dependent variable. Expected returns are estimated using the Fama-French benchmark. Robust standard errors are reported in parentheses.

*, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively. Columns (1)-(4) represent different model specifications.

Table 4

Cross-Sectional Regression Results for CAR[0, 20]

Variables	(1)	(2)	(3)	(4)
External_Revenue_AI	—	—	0.072** (0.033)	0.077** (0.033)
External_Revenue_AI × ROA	—	—	—	-0.008 (0.008)
Ext_Application	0.079** (0.032)	0.080** (0.031)	—	—
Comp_Stage	0.124*** (0.043)	0.146*** (0.049)	0.121*** (0.045)	0.124*** (0.045)
Comp_Stage × SmallCap	—	-0.106 (0.101)	—	—
SmallCap	-0.086** (0.036)	0.009 (0.097)	-0.087** (0.037)	-0.096** (0.038)
LargeCap	-0.042 (0.045)	-0.054 (0.044)	-0.046 (0.045)	-0.049 (0.045)
ROA	-0.007* (0.004)	-0.007 (0.004)	-0.008* (0.004)	-0.002 (0.007)
Leverage	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)	-0.001* (0.001)
PB	-0.003 (0.004)	-0.003 (0.004)	-0.002 (0.004)	-0.001 (0.004)
Liquidity	0.000* (0.000)	0.000* (0.000)	-0.000 (0.000)	-0.000 (0.000)
Volatility	-0.161* (0.083)	-0.170** (0.081)	-0.177** (0.084)	-0.177** (0.084)
Constant	0.035 (0.089)	0.049 (0.071)	0.098 (0.072)	0.099 (0.072)

Note: This table reports cross-sectional regression results using cumulative abnormal returns over the [0, 20] event window as the dependent variable. Expected returns are estimated using the Fama-French benchmark. Robust standard errors are reported in parentheses.

*, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively. Columns (1)-(4) represent different model specifications.

This suggests that smaller firms tend to receive less favorable abnormal returns following AI adoption announcements, potentially reflecting investor concerns about execution capacity, commercialization ability, or the credibility of implementation. By contrast, the *LargeCap* indicator is only weakly significant in the short-term window, suggesting that the market penalty is more concentrated among smaller firms rather than reflecting a simple monotonic size effect.

In the medium-term window [0, 20], the importance of firm risk and implementation maturity becomes more pronounced. As shown in Table 4, the variable *Volatility* is significantly negative only in the longer window, with an estimated effect between approximately -18% and -16%, while remaining insignificant in the short-term specifications. This result suggests that although the market may initially respond positively or neutrally to AI-related announcements, firms with higher pre-event stock volatility tend to experience weaker medium-term abnormal performance. This pattern is consistent with the interpretation that the market becomes more selective over time, discounting firms perceived as riskier or less likely to sustain the initial announcement effect.

Similarly, *comp_stage* is significantly positive only in the [0, 20] window across the medium-term specifications reported in Table 4, with an estimated effect of approximately 11% to 14%. This suggests that investors place increasing value on announcements that reflect a more advanced implementation stage as time passes. While early market reactions may be driven partly by narrative or attention effects, the medium-term response appears to reward firms that provide stronger signals of actual execution and implementation credibility.

Taken together, these results suggest that the stock market reaction to AI adoption announcements is shaped by both the economic content of the disclosure and the underlying characteristics of the firm. In the short run, investors appear to reward announcements that imply stronger revenue potential or internal technological capability. Over a longer horizon, however, the market becomes more sensitive to whether the disclosed AI strategy appears credible, executable, and sustainable, while penalizing firms with higher perceived risk or weaker implementation capacity.

As an additional robustness check, the regressions further control for pre-event AI-sector performance measured over the [-5, -1] trading-day window prior to the announcement date. This specification is intended to account for the possibility that the observed abnormal returns merely reflect pre-existing AI-sector momentum or information leakage before the formal disclosure. The coefficient on the AI-sector return variable (ai_lag) is not statistically significant in either the [0, 5] or [0, 20] specifications, while the main announcement-related variables remain qualitatively similar to the baseline results. In particular, ext_application remains positively significant across both event windows. These findings suggest that the observed abnormal returns are not primarily driven by pre-event information leakage or broad AI-sector movements, but instead reflect firm-specific market reactions to the content of AI adoption announcements. Detailed results are reported in Appendix Table A3.

5 Conclusion

This paper examines how the stock market reacts to AI adoption announcements made by Chinese A-share listed firms. Using a hand-verified sample of firm-level AI-related disclosures from February to August 2025, this study combines large language model (LLM)-based textual classification with an event study framework and cross-sectional regression analysis to evaluate both the timing and heterogeneity of investor responses.

The results suggest that AI adoption announcements do generate economically meaningful abnormal stock returns, but the market reaction is far from uniform. At the aggregate level, abnormal returns tend to rise following the announcement date, although the effect varies across event windows and firm types. The distributional evidence further shows that market responses are heterogeneous rather than universally positive, indicating that investors differentiate across firms and announcement content rather than simply reacting to AI-related keywords.

The cross-sectional results provide more detailed evidence on what drives these differences. In the short run, investors appear to reward announcements that imply stronger commercial potential, particularly those related to external value creation and revenue-relevant AI applications. Firms that demonstrate internal AI development capability also receive a stronger immediate market response, suggesting that the market values not only AI adoption itself but also the perceived strategic depth behind it. In the medium term, however, the market becomes more selective. Firms with more advanced implementation signals tend to experience stronger abnormal returns, while firms with higher pre-event stock volatility or smaller market capitalization tend to be discounted more heavily.

These findings contribute to the growing literature on market reactions to corporate technology disclosures by focusing specifically on AI adoption announcements, a category of firm-level information that has become increasingly important but remains underexplored in empirical finance research. More broadly, this paper also illustrates how LLM-based classification methods can be integrated into financial event studies to extract structured economic signals from unstructured corporate disclosures. In this sense, the study contributes not only to the understanding of AI-related market behavior but also to emerging methodological approaches in finance research.

At the same time, several limitations should be acknowledged. First, the sample period is relatively short and concentrated in a specific stage of heightened AI market attention, which may limit the generalizability of the results across time. Second, although the classification pipeline combines multiple LLM outputs with manual review, some degree of subjectivity in text classification may remain. Third, the analysis focuses primarily on short-to-medium term abnormal returns, and therefore does not fully address whether the market reaction to AI adoption announcements is ultimately sustained or reversed over longer horizons.

Future research could extend this study in several directions. One promising avenue would be to examine buy-and-hold abnormal returns (BHAR) over longer post-event horizons in order to test whether the market initially overreacts to AI-related narratives and later corrects. Another extension would be to investigate industry spillovers or competitive effects, especially in settings where leading firms' AI adoption may affect the valuation of smaller peers.

Overall, this paper suggests that financial markets do not simply reward the mention of AI. Rather, investors appear to respond to whether the announced AI strategy is commercially relevant, credibly implemented, and supported by firm-level capability. In that sense, the market reaction to AI adoption announcements reflects not only enthusiasm for a new technology, but also a more selective process of evaluating which firms are most likely to translate AI adoption into real economic value.

Reference

Chan, L. K. C., Lakonishok, J., & Sougiannis, T. (2001). The stock market valuation of research and development expenditures. *The Journal of Finance, 56*(6), 2431–2456. https://doi.org/10.1111/0022-1082.00411

Hirshleifer, D., Hsu, P.-H., & Li, D. (2013). Innovative efficiency and stock returns. *Journal of Financial Economics, 107*(3), 632–654. https://doi.org/10.1016/j.jfineco.2012.09.011

Jegadeesh, N., & Wu, D. (2013). Word power: A new approach for content analysis. *Journal of Financial Economics, 110*(3), 712–729. https://doi.org/10.1016/j.jfineco.2013.08.018

Kothari, S. P., & Warner, J. B. (2007). Econometrics of event studies. In B. E. Eckbo (Ed.), *Handbook of corporate finance: Empirical corporate finance* (Vol. 1, pp. 3–36). Elsevier. https://doi.org/10.1016/B978-0-444-53265-7.50015-9

Loughran, T., & McDonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *The Journal of Finance, 66*(1), 35–65. https://doi.org/10.1111/j.1540-6261.2010.01625.x

MacKinlay, A. C. (1997). *Event Studies in Economics and Finance.* Journal of Economic Literature, 35(1), 13–39. http://www.jstor.org/stable/2729691

Jegadeesh, N., & Wu, D. (2013). Word power: A new approach for content analysis. Journal of Financial and Quantitative Analysis, 48(5), 1243–1272. https://www.sciencedirect.com/science/article/pii/S0304405X13002328

Takeda, F., Takeda, K., Takemura, T., & Ueda, R. (2021). The impact of information technology investment announcements on the market value of the Japanese regional banks. Finance Research Letters, 41, 101811. https://www.sciencedirect.com/science/article/pii/S1544612320316251

Tetlock, P. C. (2007). Giving content to investor sentiment: The role of media in the stock market. *The Journal of Finance, 62*(3), 1139–1168. https://doi.org/10.1111/j.1540-6261.2007.01232.x

Tetlock, P. C., Saar-Tsechansky, M., & Macskassy, S. (2008). More than words: Quantifying language to measure firms' fundamentals. *The Journal of Finance*, 63(3), 1437–1467.

Appendix

Table A1

Decomposition of External-Facing AI Application Categories

Variables	[0, 5]	[0, 20]
Product_AI	0.040(0.032)	0.034(0.037)
Service_AI	0.006(0.062)	-0.001(0.071)
Industry_Solution	0.121***(0.034)	0.104***(0.039)
Comp_Stage	0.051(0.039)	0.114**(0.045)
SmallCap	-0.072**(0.033)	-0.086**(0.037)
LargeCap	-0.048(0.041)	-0.031(0.046)
ROA	-0.002(0.004)	-0.006(0.004)
Leverage	-0.000(0.001)	-0.001*(0.001)
PB	-0.002(0.004)	-0.002(0.004)
Liquidity	-0.000(0.000)	-0.000(0.000)
Volatility	0.021(0.074)	-0.163*(0.084)
Constant	0.001(0.063)	0.096(0.072)

Note: This table reports decomposition regressions that replace the aggregate *external_revenue_ai* indicator with more detailed external-facing AI categories, including *product_ai*, *service_ai*, and *industry_solution*. The dependent variable is cumulative abnormal return (CAR) over the [0, 5] and [0, 20] event windows, respectively. Expected returns are estimated using the Fama–French three-factor benchmark. Robust standard errors are reported in parentheses.

*, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table A2

Subsample Regression Results for Firms with External AI Adoption

Variables	**[0, 5]**	**[0, 20]**
Internal_AI_Model	0.056**(0.028)	0.053*(0.031)
Ext_Application	0.082***(0.031)	0.078**(0.035)
Comp_Stage	0.059(0.040)	0.123***(0.045)
SmallCap	-0.074**(0.034)	-0.086**(0.038)
LargeCap	-0.072*(0.043)	-0.046(0.049)
ROA	-0.003(0.004)	-0.007(0.004)
Leverage	0.000(0.001)	-0.001(0.001)
PB	-0.004(0.004)	-0.003(0.004)
Liquidity	-0.000(0.000)	-0.000*(0.000)
Volatility	0.035(0.082)	-0.151(0.092)
Constant	-0.028(0.066)	0.065(0.074)

Note: This table reports subsample regression results for firms that had already adopted external AI technologies. The variable intern*al_ai_model* captures whether the firm also demonstrates internal AI development capability or a mixed strategy combining internal and external AI adoption. The dependent variable is cumulative abnormal return (CAR) over the [0, 5] and [0, 20] event windows, respectively. Expected returns are estimated using the Fama–French three-factor benchmark. Robust standard errors are reported in parentheses.

*, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table A3

Robustness Check: Controlling for Pre-Event AI-Sector Performance

Variables	[0, 5]	[0, 20]
ROA	-0.003 (0.004)	-0.007* (0.004)
Leverage	0.000 (0.001)	-0.001 (0.001)
PB	-0.004 (0.004)	-0.003 (0.004)
Liquidity	-0.000 (0.000)	-0.000* (0.000)
Volatility	0.012 (0.073)	-0.163* (0.084)
AI_Lag	0.017 (0.011)	0.006 (0.013)
Ext_AI	0.014 (0.047)	0.031 (0.054)
Mix_AI	0.067 (0.048)	0.083 (0.055)
Ext_Application	0.073*** (0.028)	0.077** (0.032)
Comp_Stage	0.063 (0.038)	0.125*** (0.044)
SmallCap	-0.075** (0.032)	-0.087** (0.037)
LargeCap	-0.074* (0.040)	-0.044 (0.045)

Note: This table reports robustness regressions controlling for pre-event AI-sector performance measured over the [-5,-1] trading-day window prior to the announcement date (*AI_Lag*). The dependent variables are cumulative abnormal returns over the [0,5] and [0,20] event windows estimated using the Fama–French three-factor benchmark. Robust standard errors are reported in parentheses.

*, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.