When LLMs Go Abroad: Foreign Bias in AI Financial Predictions

Sean Cao
Robert H. Smith School of Business
University of Maryland

Charles C.Y. Wang
Harvard Business School and ECGI

Xiang Yi*
Hong Kong Polytechnic University

September 2025

Abstract

We document foreign biases in AI-generated financial predictions: ChatGPT (US-based) is systematically more optimistic about Chinese firms than DeepSeek (China-based), predicting higher end-of-year stock prices and generating more buy recommendations. This AI-specific phenomenon contradicts the traditional home bias in which investors favor domestic assets. We trace this bias to differential information access: ChatGPT's optimism increases when US media coverage of Chinese firms' negative news is scarce relative to Chinese media. Supporting this mechanism, placebo tests with synthetic Chinese firms without such asymmetries show no prediction gap between models. Crucially, providing ChatGPT with Chinese news through prompts—which cannot alter model weights—completely eliminates the prediction gap, demonstrating that the bias stems from missing training data. Our findings imply that the parallel development of LLMs in different countries can create divergent financial forecasts, potentially amplifying rather than reducing cross-border information asymmetries as these tools shape investment decisions globally.

Keywords: Artificial intelligence; Financial statement analysis; Large language models; ChatGPT; DeepSeek

JEL: D83, L86, G14, G15, G41, M41, O33

*Cao (scao824@umd.edu) is an associate professor and the director and co-founder of the AI Initiative for Capital Market Research at the Robert H. Smith School of Business. Wang (charles.cy.wang@hbs.edu) is the Tandon Family Professor of Business Administration at Harvard Business School and a Research Member of the European Corporate Governance Institute (ECGI). Yi (yi.xiang@polyu.edu.hk) is an assistant professor of accounting at Hong Kong Polytechnic University. Sean Cao acknowledges support from the AI Initiative for Capital Markets Research at the University of Maryland. The authors used Claude.ai to assist with wording and phrasing. The intellectual content, analysis, and conclusions are solely those of the authors, who remain responsible for any errors. Comments are welcome.

1 Introduction

The rapid adoption of Large Language Models (LLMs) in financial analysis represents a fundamental shift in how investment information is produced and consumed (Engelberg, Manela, Mullins, and Vulicevic, 2025; He, Lv, Manela, and Wu, 2025). ChatGPT reached 100 million users within two months of its November 2022 launch, with surveys indicating that over 40% of institutional investors now use AI tools for market analysis (Bloomberg Intelligence, 2024); other market participants also appear to be integrating the tool into their workflow (Blankespoor, DeHaan, and Li, 2025; Cheng, Lin, and Zhao, 2025; Law and Shen, 2025). A nascent literature documents LLMs' ability to process or synthesize financial statements (Bertomeu, Lin, Liu, and Ni, 2025; Bradshaw, Ma, Yost, and Zou, 2025; de Kok, 2025; Shaffer and Wang, 2024), and forecast earnings (Cook, Kazinnik, Hansen, and McAdam, 2023; Kim, Muhn, and Nikolaev, 2024; Kim and Nikolaev, 2025), corporate policies (Jha, Qian, Weber, and Yang, 2024), and stock returns (Lopez-Lira and Tang, 2023; Chen, Green, Gulen, and Zhou, 2024).

However, a growing body of work in finance reveals potential perils in applying LLMs to financial analysis. Recent studies document systematic biases including extrapolation errors and miscalibration (Chen et al., 2024), look-ahead bias and training leakage (Levy, 2025; Sarkar and Vafa, 2024; He et al., 2025), and other behavioral anomalies that can distort predictions (Bini, Cong, Huang, and Lawrence J, 2025). These biases arise from both training corpora composition and modeling choices. These documented biases raise fundamental questions about how such issues might manifest when different countries develop their own LLMs with distinct training approaches and data sources.

The parallel development of frontier models in the US and China provides a unique opportunity to address this gap, particularly critical as AI tools increasingly mediate cross-border capital flows where information asymmetries are most severe. Cross-border model comparisons offer three key advantages. First, studying a single model in isolation makes it difficult to define a natural benchmark for what constitutes bias. By contrast, when two frontier models trained in different

environments analyze identical firms, their systematic prediction differences provide a comparative benchmark that helps distinguish between model-specific errors and broader patterns in how AI systems process financial information. Second, comparing two models' predictions for identical firms helps probe potential sources of bias, which would be difficult to disentangle from prediction errors of a single model. Third, as the United States and China emerge as the two leading AI powers, understanding how their leading models differ has direct implications for both practitioners and academics. These differences are likely to shape cross-border investment analysis, market efficiency, and the global diffusion of AI.

Our investigation draws parallels between human and algorithmic decision-making under information asymmetries. While humans exhibit home bias—overweighting domestic or local assets despite diversification benefits (French and Poterba, 1991; Cooper and Kaplanis, 1994; Tesar and Werner, 1995; Lewis, 1999; Coval and Moskowitz, 1999)—the mechanism remains debated: information advantages (Portes and Rey, 2005; Van Nieuwerburgh and Veldkamp, 2009) versus behavioral biases that lead investors to favor familiar assets (Heath and Tversky, 1991; Huberman, 2001). The question of whether LLMs exhibit similar patterns is compelling because they learn from human-generated text that could embody human biases (Blodgett, Barocas, Daumé, and Wallach, 2020; Bender, Gebru, McMillan-Major, and Shmitchell, 2021; Shumailov, Shumaylov, Zhao, Gal, Papernot, and Anderson, 2023), or could produce new biases due to how information incorporated. This creates two possibilities for cross-border bias in LLMs: foreign information might be absent from training (data availability) or present but underweighted (data weighting). These may appear observationally equivalent—whether an LLM lacks foreign information or learned to ignore it, the result is systematic prediction differences we can measure and manipulate.

This paper documents a striking reversal of traditional home bias when LLMs analyze foreign firms. While humans favor familiar domestic assets, we find ChatGPT (US-based) is systematically more optimistic than DeepSeek (China-based) about Chinese firms, exhibiting what we term "foreign bias." We trace this to asymmetric information processing.

We prompt both ChatGPT 4.1 and DeepSeek R1—frontier models from their respective countries—to provide six-month price predictions and buy/hold/sell recommendations for Chinese firms as of June 30, 2024. This design allows us to observe how models trained on different corpora assess identical firms at the same point in time. To establish causality, we supplement ChatGPT's information set with Chinese news articles before requesting predictions, effectively equalizing information access across models.

Predictions generated by ChatGPT are 2.415 RMB (11.6%) higher than DeepSeek and recommendations generated by ChatGPT are 1.3 percentage points more likely to be "buy," despite DeepSeek having superior effective access to Chinese-language information about these firms. This foreign bias is not explained by firm characteristics—the differences persist after controlling for size, profitability, leverage, past returns, and other observables that might influence AI assessments. Crucially, no such differences exist when both models analyze US firms, confirming the bias is specific to cross-border predictions rather than reflecting general model differences.

We trace this foreign bias to differential information access between models. Using comprehensive news data from RavenPack (US media) and the China Finance News Database (Chinese media), we construct firm-level measures of coverage gaps. Chinese media provides 87% of negative news and 93% of positive news about Chinese firms, with many firms receiving zero US coverage. ChatGPT's excess optimism increases precisely when US media coverage is scarce, the interaction between the ChatGPT indicator and news gap measures is economically and statistically significant. When US sources provide little information, ChatGPT appears to fill the void with optimistic priors rather than tempering its predictions.

Our strongest evidence comes from directly manipulating ChatGPT's information set. When we provide ChatGPT with Chinese news articles about each firm before requesting predictions, its excess optimism completely disappears. This intervention is powerful precisely because it distinguishes between the data availability and data weighting mechanisms: prompting can only provide new information without changing the underlying model weights. Since this information provision

alone eliminates the bias—the coefficient on the ChatGPT indicator becomes -0.798 (insignificant) for prices and -0.002 (insignificant) for recommendations, compared to baseline differences of 2.415 and 0.013 respectively—we can attribute the foreign bias to data availability. In other words, ChatGPT's relative optimism seems to stem from lacking Chinese news in its training corpus, rather than from having but underweighting such information.

In addition, ChatGPT's relative optimism represents bias rather than superior information. We show that ChatGPT's predictions are significantly less accurate than DeepSeek's: it exhibits 13.3% larger absolute price prediction errors and 7% higher recommendation error rates when compared to realized December 31, 2024 outcomes. The pattern of errors—largest for high-leverage and high-volatility firms—suggests ChatGPT struggles most when Chinese-specific information is likely crucial for valuation, further confirming that missing local data, rather than processing differences, drives the foreign bias.

Robustness tests confirm our findings across multiple dimensions. The foreign bias persists for twelve-month predictions, ruling out horizon-specific effects. Enabling web search functionality reduces, but does not eliminate, the bias (coefficient drops 78% but remains significant), suggesting real-time information access only partially compensates for missing training data. Placebo tests using synthetic Chinese firms show no prediction differences, confirming the bias requires actual information asymmetry rather than reflecting mechanical model differences.

This paper makes three main contributions to the literature. First, we document a novel form of bias in AI systems that reverses traditional behavioral patterns. While humans exhibit home bias by favoring domestic assets they know better, AIs exhibit foreign bias by being overly optimistic about foreign assets they know less about. This reversal highlights how algorithmic decision-making, even when trained on human-generated data, can produce behavioral phenomena distinct from human psychology when operating under different information constraints.

Second, we expand understanding of AI bias origins. While existing literature often focuses on endogenous sources of bias—intentional modeling choices by developers or unintentional human

errors in the development process——our cross-border comparison reveals that AI bias can also arise from exogenous differences in informational environments. Our evidence indicates that the systematic prediction differences between ChatGPT and DeepSeek stem from the fundamentally different information landscapes in which these models were trained. ChatGPT's training corpus naturally contains more English-language sources, while DeepSeek's contains more Chinese sources, creating differential information access that manifests as systematic bias. Our information intervention provides strong evidence for this mechanism: supplying ChatGPT with Chinese news through prompts, which cannot alter model weights, eliminates the bias, pointing to data availability rather than data weighting as the primary driver. This finding suggests that even with perfect development practices, AI systems will inherit biases from the information environments they inhabit, making cross-border AI applications inherently susceptible to foreign bias.

Third, we expand the information asymmetry literature by analyzing a fundamentally different type of information intermediary. Unlike analysts or media, who can actively seek out information to overcome asymmetries (Brennan and Subrahmanyam, 1995; Bushee, Core, Guay, and Hamm, 2010), LLMs are constrained by their training data and learned weights, which can codify or even amplify information gaps. By studying the two leading frontier models—ChatGPT in the United States and DeepSeek in China—we show how models trained in different contexts can yield systematically different assessments of the same firms. As the U.S. and China emerge as dominant AI powers, these divergences have important implications for practitioners, policymakers, and scholars concerned with cross-border investment, market efficiency, and the global diffusion of AI technologies.

Our findings have important practical implications. For practitioners, the results highlight the importance of understanding AI tool limitations in financial analyses (Cao, Jiang, Wang, and Yang, 2024; Levy, 2025; Lee, Seo, Park, Lee, Ahn, Choi, Lopez-Lira, and Lee, 2025), particularly when making cross-border investment decisions. The fact that ChatGPT is systematically over-optimistic about Chinese firms—and potentially other foreign markets with limited English-

language coverage—suggests investors should not rely solely on these tools for international analysis. The ease with which we eliminate the bias by providing additional information suggests a simple solution: supplementing AI analysis with local information sources. Investment firms might benefit from using multiple AI models trained in different languages or explicitly fine-tuning models on balanced international data.

For researchers, our paper demonstrates the importance of studying AI behavior as distinct from human behavior. While LLMs are trained to mimic human language and reasoning, their behavioral patterns can diverge significantly from human psychology when operating under different constraints. The reversal of home bias shows that behavioral finance insights do not automatically transfer to algorithmic decision-makers. As AI tools become more prevalent in financial markets, understanding their specific biases and limitations becomes essential for predicting market outcomes and designing appropriate regulations.

Looking forward, our results raise questions about the market-wide effects of AI adoption. If many investors simultaneously rely on biased AI tools, systematic mispricing could emerge, potentially creating profitable opportunities for sophisticated arbitrageurs aware of these biases. The growing influence of AI on cross-border capital flows may reshape global financial markets in unexpected ways, potentially reversing decades-long trends in international portfolio allocation. As the AI ecosystem evolves with new models and training approaches, continuous monitoring of these biases will be essential for maintaining market efficiency and protecting investors.

2 Sample Construction and Data

2.1 Sample Selection

Our sample construction begins with the universe of Chinese firms listed on the Shanghai and Shenzhen stock exchanges with available data in Compustat Global. We focus on Chinese firms as they represent a major cross-border investment opportunity where information asymmetries

between US and Chinese investors are particularly pronounced. The parallel development of sophisticated LLMs in both countries—ChatGPT in the US and DeepSeek in China—provides an ideal setting to study how AI models trained in different linguistic and informational environments assess the same set of firms.

We start with all Chinese firms in Compustat Global as of December 31, 2023, requiring non-missing data for key financial variables used as controls in our analysis. Specifically, we require firms to have available information on total assets (to calculate $FIRMSIZE$), return on assets (ROA), and leverage (LEV). These screens ensure that both AI models have access to basic financial information when making predictions, allowing us to isolate the effect of information asymmetries in news and contextual information rather than differences in fundamental financial data availability. This process yields 4,978 unique Chinese firms.

2.2 AI Prediction Collection

For each firm in our sample, we collect stock price predictions and investment recommendations from two frontier LLMs: ChatGPT 4.1 (developed by OpenAI) and DeepSeek R1 (developed by DeepSeek). We access both models through their respective APIs to ensure consistency and replicability. It should be noted that both ChatGPT4.1 and DeepSeek R1's training sample periods ended on June 30, 2024, with all predictions made for the same future date of December 31, 2024, creating a six-month out-of-sample prediction horizon. This can mitigate the influence of the forward-looking bias of LLM on our results (Didisheim, Fraschini, and Somoza, 2025).

We use standardized prompts to ensure comparability across models. For price predictions and buy recommendation, we prompt each model with: "Here is a firm called "name" (ticker: ticker), listed as an A-share in China. As a professional financial analyst, please: 1. Predict its stock price on 31 December 2024 (three decimal places). 2. Predict whether the price will increase from the end of June this year to December 31, 2024 (1 = up, 0 = not up). Do NOT use current online data or random numbers - use your best effort." We randomize the order in which firms are presented

to each model to avoid any potential ordering effects. To ensure data replicability, we set the temperature to zero. To ensure data quality, we re-query both models for all our sample firms one week later to verify consistency, finding correlation coefficients exceeding 0.95 for both models' predictions.

This process generates 9,956 firm-AI observations (4,978 firms $\times$ 2 AIs), with each observation containing a price prediction (*AI_PREDICTION*) and a buy recommendation indicator (*AI_RECOMMEND*, coded as 1 for buy recommendations and 0 for hold or sell recommendations).

2.3 News Coverage Data

To examine the mechanism driving AI prediction differences, we collect comprehensive data on media coverage from both US and Chinese sources. This allows us to construct differences in the training sample between the two AI models. Earlier, we mentioned that ChatGPT's training prioritizes US sources, whereas DeepSeek gives more importance to local Chinese news.

For US media coverage, we use RavenPack News Analytics, which processes news from major US financial publications including the Wall Street Journal, Financial Times, Bloomberg, Reuters, and Dow Jones Newswires (Drake, Guest, and Twedt, 2014; Guest, 2021; Lock, 2024). RavenPack provides sentiment scores for each news article mentioning our sample firms. We aggregate all news articles from June 30, 2021, to June 30, 2024 (a three-year window) to capture the information environment likely reflected in the LLMs' training data. Articles are classified as negative if they have a sentiment score below 0 on RavenPack's [-1, +1] scale, and positive if above 0.

For Chinese media coverage, we use the China Finance News Database (CFND), which aggregates content from major Chinese financial media outlets including Caixin, Securities Times, China Securities Journal, and Shanghai Securities News (Li, Li, Wang, and Thatcher, 2021; Gao, Zhang, and Yang, 2023; Qiao, 2023). Similar to our approach with RavenPack, we collect all articles mentioning our sample firms during the same three-year period and classify them by sentiment using CFND's proprietary sentiment analysis, which is calibrated to be comparable to RavenPack's

methodology.

We then construct two firm-level measures of cross-border information gaps. *NEG_NEWS_GAP* is defined as

$$\frac{\text{CN negative} - \text{US negative}}{\text{CN negative} + \text{US negative}}$$

computed over the three-year window prior to the models' training cutoff (June 30, 2021 to June 30, 2024). Here, *CN negative* refers to the number of negative news articles about a given firm from Chinese media sources (CFND), while *US negative* refers to the number of negative news articles from U.S. media sources (RavenPack). This variable lies in $[-1, 1]$, with $+1$ indicating only Chinese negative coverage and -1 indicating only U.S. negative coverage.

NET_NEG_NEWS_GAP is defined as

$$\frac{(CN\ negative - CN\ positive) - (US\ negative - US\ positive)}{CN\ negative + CN\ positive + US\ negative + US\ positive}$$

where *CN positive* (*US positive*) denotes the number of positive news articles about the firm from Chinese (U.S.) media sources. This measure captures the net negativity gap across countries, with higher values indicating that Chinese outlets are, on net, more negative than U.S. outlets. For firms with zero articles in both markets (denominator $= 0$), we set the gap to zero; results are robust to dropping such firms or adding a no-news indicator.

2.4 Control Variables

We obtain firm-level financial data from Compustat Global measured as of fiscal year-end 2023, ensuring all control variables precede our AI predictions. Our control variables include standard predictors of stock returns and firm value from the asset pricing literature. *FIRMSIZE* is the natural logarithm of total assets. *LEV* is total debt divided by total assets. *ROA* is income before extraordinary items divided by total assets. *LOSS* is an indicator variable equal to one if the firm reported negative net income. *BM* is the book-to-market ratio, calculated as book

value of equity divided by market value of equity. *PRE_RETURN* is the stock return during 2023. *RETURN_VOLATILITY* is the standard deviation of daily returns during 2023. *IO* is the percentage of shares held by institutional investors, obtained from the China Stock Market and Accounting Research (CSMAR) database.

For robustness tests and mechanism analyses, we collect additional data. Foreign institutional ownership (*FOREIGN_IO*) comes from CSMAR's detailed ownership files. Actual stock prices on December 31, 2024, used to calculate prediction errors, are obtained from Compustat Global.

2.5 Sample Composition and Descriptive Statistics

Table 1, Panel A provides descriptive statistics for our key variables. The AI prediction variables reveal interesting baseline patterns. *AI_PREDICTION* has a mean of 20.76 RMB with substantial variation (standard deviation = 16.07). The distribution is right-skewed with median (15.33) below mean, and interquartile range spanning 8.32 to 28.46 RMB. Strikingly, *AI_RECOMMEND* has a mean of 0.98, indicating 98% of all AI assessments are buy recommendations. This extreme optimism from both models provides a high baseline against which to measure differential optimism.

The information gap measures confirm severe asymmetries. The mean *NEG_NEWS_GAP* of 0.58 and mean *NET_NEG_NEWS_GAP* of 0.05 show that Chinese media tend to report more negative corporate news than US media. This suggests that ChatGPT 4.1 tends to omit many negative media news from Chinese companies in its training sample.

Firm characteristics show reasonable variation for regression analysis. Average log assets of 8.598 represent approximately 5.4 billion RMB. Leverage averages 42%, ROA averages 2%, and 21% of firm-years show losses. The mean book-to-market ratio of 0.57 suggests moderate growth expectations, while prior returns average 3%. Institutional ownership averages 41.1%, ranging from minimal to over 60% at the 75th percentile.

Table 1, Panel B presents the industry distribution of our sample using the Fama-French 12-industry classification. Manufacturing comprises the largest sector (24.89% of observations), fol-

lowed by Business Equipment (20.37%), reflecting the industrial composition of Chinese capital markets. Healthcare (8.44%) and Chemicals (8.86%) represent other significant sectors, while Financial Services accounts for 4.38% of the sample. In addition, the measures of the negative news gap reveal striking disparities in media coverage across borders. This suggests the importance of controlling for industry fixed effect in our empirical analysis.

3 Empirical Results

3.1 Foreign Bias in AI-Generated Financial Predictions

Table 2 presents our primary results examining systematic differences in financial predictions between ChatGPT and DeepSeek for Chinese firms. The dependent variable in Column (1) is *AI_PREDICTION*, representing each AI's predicted stock price for December 31, 2024, approximately six months after the training sample end date of June 30, 2024. Our key independent variable is *US_AI*, an indicator variable equal to one for predictions generated by ChatGPT 4.1 and zero for those from DeepSeek R1.

Column (1) reveals that ChatGPT predicts systematically higher stock prices than DeepSeek. The coefficient on *US_AI* is 2.4153, indicating that ChatGPT's price predictions exceed DeepSeek's by approximately 2.4 RMB per share. This difference is economically substantial, representing 11.6% of the mean predicted price and approximately 15.7% of the median prediction. To contextualize this magnitude, a 2.4 RMB difference for the median firm translates to a market capitalization difference of approximately 240 million RMB, assuming 100 million shares outstanding.

The effect persists after controlling for a comprehensive set of firm characteristics that might influence AI predictions. We control for *FIRMSIZE* (natural logarithm of total assets), *LEV* (total debt divided by total assets), *ROA* (income before extraordinary items divided by total assets), *LOSS* (indicator for negative net income), *BM* (book value of equity divided by market value of equity), *PRE_RETURN* (stock return in the prior year), *RETURN_VOLATILITY* (standard

deviation of daily returns in the prior year), and *IO* (percentage of shares held by institutional investors). All control variables are measured as of year-end 2023 to ensure they precede the AI predictions. We also include industry fixed effects based on Fama-French 12-industry classifications to account for systematic sectoral differences.

The control variables exhibit intuitive relationships with price predictions. More profitable firms (*ROA*) receive higher price predictions (33.7108, $p < 0.01$), suggesting AIs view highly profitable firms as a value of the upside potential. Firms with higher book-to-market ratios (*BM*) also receive lower predictions (coefficient = -4.1948, $p < 0.01$), suggesting that AIs favor growth characteristics. Prior returns (*PRE_RETURN*) strongly negatively predict future price targets (coefficient = -8.6622, $p < 0.01$), suggesting the AI models incorporate momentum reversal or mean reversion in their forecasting approaches.

Column (2) examines buy recommendations, where *AI_RECOMMEND* equals one for buy recommendations and zero for hold or sell recommendations. The coefficient on *US_AI* is 0.0135 ($p < 0.01$), indicating ChatGPT issues 1.35 percentage points more buy recommendations than DeepSeek. While this difference appears modest in absolute terms, it is substantial given the high baseline recommendation rate. With DeepSeek recommending buys for 97.55% of firms, ChatGPT's 98.85% buy rate represents a 53% reduction in non-buy recommendations (from 2.45% to 1.15%). This near-universal optimism about Chinese equities from both AIs is itself noteworthy, but ChatGPT's even greater enthusiasm supports our foreign bias hypothesis.

The adjusted R-squared values differ markedly between specifications—0.307 for price predictions versus 0.033 for recommendations—suggesting price targets incorporate firm characteristics more systematically than binary recommendations. The low R-squared for recommendations partly reflects the limited variation in the dependent variable but also indicates that the AI recommendation differences are not driven by observable firm characteristics.

3.2 Information Asymmetry as the Underlying Mechanism

Table 2 examines whether cross-border differences in media coverage explain ChatGPT's relative optimism about Chinese firms. Specifically, we interact the U.S. AI indicator (*US_AI*) with two measures of media asymmetry: the relative gap in negative news coverage (*NEG_NEWS_GAP*) and the net difference in negative news coverage (*NET_NEG_NEWS_GAP*). These variables capture the extent to which U.S. versus Chinese media report negative information about Chinese firms. If ChatGPT systematically underweights Chinese-language news, we would expect its optimism to increase precisely when U.S. sources provide less negative coverage.

Columns (1)–(2) present results for AI price predictions. In Column (1), the interaction between *US_AI* and *NEG_NEWS_GAP* is strongly positive (coefficient = 3.5373, $p < 0.01$), while the main effect of *NEG_NEWS_GAP* is negative and significant. This implies that when U.S. outlets report relatively fewer negative stories than Chinese outlets, ChatGPT becomes more optimistic relative to DeepSeek, consistent with information underweighting. In Column (2), results are even more pronounced: the interaction between *US_AI* and *NET_NEG_NEWS_GAP* is large (30.2682, p < 0.01), confirming that U.S.–China disparities in negative coverage amplify ChatGPT's upward bias in stock price forecasts.

Columns (3)–(4) examine AI recommendations. The interaction coefficients again show positive and significant effects. In particular, the interaction of *US_AI* with *NET_NEG_NEWS_GAP* is 0.1420 ($p < 0.01$), suggesting that ChatGPT is more likely than DeepSeek to issue buy recommendations when U.S. outlets underreport negative Chinese news. The main effect of *NET_NEG_NEWS_GAP* is negative (-0.1265, $p < 0.01$), consistent with more negative Chinese coverage leading to fewer buy recommendations overall.

The control variables behave as expected. Larger and more profitable firms (*FIRM_SIZE*, *ROA*) are associated with higher predictions, while firms with higher leverage or book-to-market ratios receive lower valuations and recommendations. Prior returns strongly negatively predict future price targets, consistent with mean reversion. Importantly, the results are robust to industry fixed

effects and clustering by Fama-French 12-industries.

Taken together, these findings provide direct evidence that asymmetric media coverage drives AI foreign bias. ChatGPT's relative optimism emerges not randomly but specifically when Chinese firms receive disproportionately negative coverage in Chinese outlets compared to U.S. outlets. This pattern highlights how gaps in training-relevant information environments can translate into systematic biases in LLM-generated financial forecasts.

3.3 Role of Media Coverage: When AI Has No News

Table 4 investigates whether the foreign bias persists when both U.S. and Chinese media coverage is absent, and whether placebo firms show similar patterns. Panel A restricts the sample to actual Chinese firms with no relevant U.S. or Chinese news in the past three years. Panel B reports results for a placebo set of synthetic Chinese firm names. In both cases, the coefficient on *US_AI* is small and statistically insignificant, suggesting that the bias requires genuine information asymmetry in media coverage.

Panel A shows that when firms have no media coverage in either market, the coefficient on *US_AI* is 1.6741 for price predictions and 0.000 for recommendations, neither of which is statistically significant. This indicates that ChatGPT does not display excess optimism relative to DeepSeek in the absence of cross-border news disparities. Control variables behave as expected: more profitable firms (*ROA*) are associated with higher predictions, while firms with greater *LEVERAGE* and higher prior *PRE_RETURN* receive lower valuations. *RETURN_VOLATILITY* and *IO* remain positively related to predictions.

Panel B presents the placebo test using synthetic Chinese firm names. Here, the coefficients on *US_AI* remain economically small and statistically insignificant for both outcomes, further confirming that the documented foreign bias is not an artifact of model differences. Instead, it arises only when there is genuine asymmetry in cross-border information. Interestingly, some firm controls (e.g., *ROA*, *LOSS*) load significantly in this placebo sample, but these effects do not alter

the interpretation of the placebo test.

Taken together, Table 4 provides strong evidence that ChatGPT's foreign bias is conditional on asymmetric media environments. When no differential coverage exists, or when the firms are purely placebo, the excess optimism disappears, highlighting the central role of cross-border information asymmetry in driving the AI foreign bias.

3.4 Closing the News Gap

Table 5 reports results from the information intervention in which ChatGPT was provided with Chinese news articles about each firm from June 2022 to June 2024, which covers three years corporate news prior to the training period end. This intervention directly addresses the information asymmetry by supplementing ChatGPT's input with local news. The key result is that the foreign bias completely disappears: the coefficient on *US_AI* becomes statistically indistinguishable from zero for both price predictions (Column 1) and buy recommendations (Column 2). This finding confirms that the documented bias is driven by asymmetric information rather than inherent model differences.

In Column (1), the coefficient on *US_AI* is negative (-0.7980) and statistically insignificant at conventional levels, compared to a significant positive effect in earlier tables. Similarly, in Column (2), the effect on recommendations is essentially zero (-0.0022). Thus, once ChatGPT has access to Chinese news, its predictions are no longer systematically more optimistic than those of DeepSeek. This evidence strongly supports the interpretation that ChatGPT's foreign bias is conditional on limited access to negative Chinese news coverage.

Control variables remain consistent with prior results. More profitable firms (*ROA*) are associated with higher price predictions, while firms with greater *LEVERAGE*, higher *BM*, and higher prior *PRE_RETURN* receive lower valuations. *RETURN_VOLATILITY* loads positively, while *IO* is again associated with higher predictions. These patterns mirror those in Tables 3 and 4, reinforcing the robustness of control variable effects.

Overall, the evidence in Table 5 identifies the likely mechanism driving the foreign bias. The information intervention—providing ChatGPT with Chinese news articles through the prompt—completely eliminates its excess optimism. This result is theoretically important because prompting fundamentally cannot alter a model's pre-trained weights; it can only provide new information to process. The fact that information provision *alone* removes the bias demonstrates that ChatGPT's foreign optimism stems from a data availability problem—the model simply lacks Chinese news in its training corpus—rather than a weighting problem where it has the information but systematically discounts it. This clean identification is possible precisely because LLMs' architecture separates information provision (through prompts) from weight updates (through training).

3.5 Boundary Condition - US firms (placebo test)

Table 6 examines whether the documented foreign bias extends to U.S. firms by repeating the baseline analysis for a placebo sample of 11,688 U.S. firm-year observations. The dependent variables are price predictions (Column 1) and buy recommendations (Column 2), with *US_AI* indicating ChatGPT predictions. If the bias were driven by general model differences, we would expect to observe similar patterns for U.S. firms. Instead, the coefficients on *US_AI* are economically small and statistically insignificant in both columns, confirming that the foreign bias is specific to cross-border predictions about Chinese firms.

Column (1) shows that ChatGPT's price predictions are indistinguishable from DeepSeek's for U.S. firms: the coefficient on *US_AI* is −0.1539. Column (2) reports a similar result for recommendations, with the coefficient on *US_AI* equal to −0.0472, again insignificant at conventional levels. These results stand in stark contrast to the strong positive coefficients observed for Chinese firms, underscoring that excess optimism arises only in the cross-border setting where information asymmetries are severe.

Taken together, Table 6 demonstrates that the foreign bias documented earlier does not generalize to all firms. This further fails to support the home bias view that ChatGPT would favor US

firms more than DeepSeek does. Instead, it is a boundary condition that emerges specifically for Chinese firms where U.S. media coverage is limited relative to Chinese coverage, providing further evidence that asymmetric cross-border information drives the bias in AI predictions.

3.6 Prediction Accuracy and Error Analysis

Table 6 compares the accuracy of ChatGPT and DeepSeek predictions against realized outcomes. Column (1) reports absolute price prediction errors, while Column (2) reports recommendation errors, defined as the absolute deviation from an ex-post optimal buy/hold/sell classification based on realized returns. Positive coefficients on *US_AI* indicate that ChatGPT is less accurate than DeepSeek, consistent with its excess optimism reflecting bias rather than superior information.

In Column (1), the coefficient on *US_AI* is 1.0863 ($p < 0.01$), showing that ChatGPT's price predictions exhibit significantly larger absolute errors than DeepSeek's. In Column (2), the coefficient on *US_AI* is 0.0074 ($p < 0.10$), indicating higher classification errors in buy/hold/sell recommendations. These findings confirm that the relative optimism documented earlier is associated with lower predictive accuracy.

Control variables load in intuitive directions. More profitable firms (*ROA*) are associated with larger prediction errors (24.6574, $p < 0.01$ for price errors; 0.9947, $p < 0.01$ for recommendation errors), suggesting that AIs tend to overestimate the performance of highly profitable firms. Firms with higher leverage (*LEVERAGE*) have smaller price errors (-6.2347, $p < 0.01$). Loss firms (*LOSS*) experience higher errors in both predictions (1.8352, $p < 0.05$) and recommendations (0.0532, $p < 0.01$). Higher book-to-market ratios (*BM*) reduce price errors (-2.8381, $p < 0.01$) but increase recommendation errors (0.0875, $p < 0.01$). Prior returns (*PRE_RETURN*) reduce price errors (-1.3305) but are insignificant for recommendations. Volatile firms (*RETURN_VOLATILITY*) exhibit significantly larger errors (2.6255, $p < 0.01$ for prices; 0.0431, $p < 0.01$ for recommendations). Institutional ownership (*IO*) is positively associated with errors in Column (1) (0.0462, $p < 0.05$), though the effect is negligible in magnitude.

Taken together, Table 6 shows that ChatGPT's foreign bias corresponds to systematically lower accuracy compared to DeepSeek. Rather than reflecting informational advantages, ChatGPT's optimism results in greater prediction and recommendation errors, reinforcing the interpretation of excess optimism as bias.

3.7 Robustness Tests

Table 8 presents three sets of robustness analyses. Panel A extends the prediction horizon to 12 months, Panel B enables web search before prediction, and Panel C repeats the baseline analysis using Mandarin prompts. Across all specifications, the coefficients on *US_AI* remain positive and significant, although attenuated in some cases, showing that the foreign bias is robust to alternative horizons, real-time information access, and prompt language.

Panel A: Alternative Prediction Dates. Column (1) shows that *US_AI* is significantly positive (1.8490, $p < 0.05$) for price predictions when extending the horizon to 12 months, while Column (2) reports a similarly positive effect on recommendations (0.0207, $p < 0.01$). These results indicate that the bias persists beyond the six-month baseline window. Control variables remain consistent: *ROA* (71.6423, $p < 0.01$) and *RETURN_VOLATILITY* (4.8821, $p < 0.01$) increase predictions, while *LEVERAGE* (-26.2111, $p < 0.01$) and *BM* (-10.1748, $p < 0.05$) reduce them.

Panel B: AI Prediction with Online Search. Column (1) shows that enabling web search reduces but does not eliminate the bias: the coefficient on *US_AI* falls to 0.5338 ($p < 0.05$), compared to larger baseline effects. Column (2) confirms a persistent recommendation bias (0.0426, $p < 0.05$). This suggests that real-time information access attenuates but cannot fully close the information gap embedded in model training. Controls again behave intuitively, with *LOSS* (0.9130, $p < 0.01$) increasing price errors, while *BM* reduces them (-1.3104, $p < 0.01$).

Panel C: AI Prediction with Mandarin Prompts. Column (1) reports a positive and significant coefficient on *US_AI* (1.5173, $p < 0.05$) even when predictions are solicited in Mandarin, while Column (2) shows a smaller but significant effect on recommendations (0.0135, $p < 0.10$).

These results confirm that the foreign bias is not simply an artifact of prompt language, although its magnitude is somewhat reduced. Control variables remain consistent: *ROA* (60.9911, $p < 0.05$) and *RETURN_VOLATILITY* (4.0061, $p < 0.01$) increase predictions, while *LEVERAGE* reduces them (-22.6838, $p < 0.01$).

Taken together, Table 8 demonstrates that the foreign bias documented in earlier tables is highly robust. It persists across longer prediction horizons, real-time web-enabled predictions, and prompt languages, underscoring that the mechanism is deeply rooted in asymmetric information processing rather than technical artifacts of prompting or prediction window choice.

3.8 Market reactions

Table 9 explores whether markets respond to the release of frontier LLMs by examining cumulative abnormal returns (CARs) around the announcements of ChatGPT 3.5 (Panel A) and DeepSeek R1 (Panel B). We relate CARs to measures of foreign institutional ownership and cross-border news gaps. We assume that foreign institutional investors are more likely to rely on ChatGPT before the release of DeepSeek R1. Thus, the foreign optimism bias in ChatGPT tends to increase the foreign investors' optimism towards Chinese firms, leading to a more positive market reaction to Chinese firms at the announcement date of ChatGPT 3.5. While these results are descriptive and not interpreted causally, they provide suggestive evidence that markets react in line with the hypothesized AI-induced biases.

Panel A: ChatGPT 3.5 Announcement. Columns (1)–(2) show that foreign institutional ownership (*IO_FOREIGN*) is positively associated with CARs. Specifically, the coefficient on *IO_FOREIGN* is 0.0031 ($p < 0.05$) for $CAR(-1, 1)$ and 0.0130 ($p < 0.01$) for $CAR(-3, 3)$. Moreover, interaction effects with negative news gaps are positive and significant: *IO_FOREIGN* $\times$ *NEG_NEWS_GAP_2019_22* loads at 0.0145 ($p < 0.05$) in Column (3) and 0.0236 ($p < 0.05$) in Column (4). Similarly, the interaction with *NET_NEG_NEWS_GAP_2019_22* is significant in Columns (5)–(6), with coefficients of 0.0107 ($p < 0.05$) and 0.0215 ($p < 0.10$), respectively. These results

suggest that markets respond more positively to ChatGPT's release when foreign ownership is high and when U.S. coverage understates negative Chinese news. Control variables behave predictably: larger firms (*FIRM_SIZE*) are positively associated with CARs, while higher *ROA*, *LEVERAGE*, and *BM* predict lower CARs.

Panel B: DeepSeek R1 Announcement. In contrast, Columns (1)–(2) show that *IO_FOREIGN* is negatively associated with CARs: -0.0013 ($p < 0.01$) for $CAR(-1, 1)$ and -0.0020 ($p < 0.05$) for $CAR(-3, 3)$. Interaction terms also indicate negative responses: *IO_FOREIGN* $\times$ *NEG_NEWS_GAP_2022_25* loads -0.0032 ($p < 0.05$) in Column (3) and -0.0056 ($p < 0.05$) in Column (4). These findings suggest that, unlike ChatGPT's U.S. launch, DeepSeek's Chinese launch elicited weaker or negative reactions from foreign investors, especially when cross-border negative news asymmetries were pronounced. Control variables again align with intuition: larger firms (*FIRM_SIZE*) are associated with lower CARs, while more profitable firms (*ROA*) show positive associations with CARs in several specifications.

Taken together, Table 9 provides suggestive evidence that financial markets react differently to AI announcements across countries. ChatGPT's release was met with relatively positive reactions among firms with higher foreign ownership, particularly when U.S. media underreported negative Chinese news, whereas DeepSeek's release elicited weaker or even negative market responses from foreign investors.

3.9 Summary and Implications

Our study uncovers a novel "foreign bias" in AI-generated financial predictions: U.S.-based ChatGPT is systematically more optimistic about Chinese firms than China-based DeepSeek. This phenomenon directly reverses the traditional home bias observed among human investors. The bias is economically meaningful—ChatGPT predicts prices 11.6% higher and issues 1.3 percentage points more buy recommendations than DeepSeek—yet disappears when the models analyze U.S. firms. We trace this bias to asymmetric information environments: ChatGPT's optimism intensifies

precisely when U.S. media provide limited coverage of Chinese firms, and providing ChatGPT with Chinese news fully eliminates the gap. Importantly, the bias translates into lower accuracy—ChatGPT exhibits larger prediction and recommendation errors—confirming that the optimism reflects informational blind spots rather than superior insight.

These findings carry several implications. First, they highlight how LLMs, trained on human-generated text, can develop biases distinct from human psychology. Whereas human investors exhibit home bias due to familiarity and comfort with domestic assets, AI models can display the opposite pattern—excessive optimism about foreign markets where their training data are sparse. Second, our results provide causal evidence that data availability and weighting shape AI behavior. A simple intervention—feeding missing news—eradicates the foreign bias, demonstrating that information asymmetry is the root mechanism. Third, the findings suggest that financial markets already respond to these AI-induced biases: firms with high foreign ownership experienced more positive abnormal returns following ChatGPT's release, while DeepSeek's launch was met with muted or negative reactions.

For practitioners, our evidence cautions against sole reliance on single-country AI systems in cross-border investment. Investors should recognize that widely used tools like ChatGPT may systematically mis-evaluate foreign firms when domestic news coverage is thin. Yet the solution is straightforward: supplementing models with local news inputs or employing multiple AIs trained in different information environments can substantially mitigate these biases. For researchers, our study demonstrates the importance of analyzing AI decision-making as a distinct object of inquiry. Rather than assuming AI will replicate or eliminate human biases, scholars must account for the ways training data, model design, and information asymmetries interact to generate new forms of systematic error.

In sum, AI adoption in finance does not merely automate human judgment—it introduces novel behavioral dynamics with real market consequences. The foreign bias we document underscores that LLMs are not neutral analyzers of global information. As their influence on capital allocation

grows, understanding and correcting such biases will be essential to ensure that AI contributes to, rather than undermines, market efficiency.

4 Conclusion

This paper documents a novel foreign bias in AI-generated financial predictions, a striking reversal of the traditional home bias observed in human investors. While humans favor familiar domestic assets, we find that ChatGPT is systematically more optimistic about Chinese firms than DeepSeek, despite DeepSeek having superior access to Chinese-language information. We trace this bias to information asymmetry between training corpora, with ChatGPT's excess optimism increasing when US media coverage of Chinese firms is scarce.

Our information-intervention evidence lends support to this hypothesized mechanism: providing ChatGPT with Chinese news articles completely eliminates its foreign bias, demonstrating that differential information access drives these systematic prediction differences. Our identification strategy exploits a key feature of LLM architecture: prompts can provide new information but cannot change pre-trained weights. Since providing information alone eliminates the bias, we can attribute it primarily to missing information in ChatGPT's training corpus rather than systematic underweighting of available information, with important implications for how to address such biases in practice.

Our results also speak to broader questions about AI development in an increasingly multipolar technological landscape. As different countries develop their own LLMs trained primarily on local-language corpora, we may see the emergence of multiple, systematically biased views of global markets. Rather than AI serving as a neutral, objective analyzer that overcomes human behavioral biases, we find that AI models can introduce new forms of bias that operate in ways that differ from traditional human biases. This suggests that the promise of AI to democratize access to sophisticated financial analysis may instead create new geographical disparities based on which models investors can access.

From a practical perspective, our findings offer both warnings and solutions. Investors should not rely solely on single AI models for cross-border investment decisions, as these models may harbor systematic biases exceeding traditional analyst biases. However, our information-intervention results provide a straightforward remedy: supplementing AI analysis with local information sources eliminates these biases. Investment firms might benefit from ensemble approaches using multiple models trained in different languages, or from fine-tuning models on balanced multilingual data to ensure comprehensive information processing.

For researchers, our paper demonstrates the importance of studying AI behavior as distinct from human behavior rather than assuming AI will simply replicate or eliminate human biases. The reversal of home bias we document suggests that insights from behavioral finance do not automatically transfer to algorithmic decision-makers. As AI tools become more prevalent in financial markets, we need to deepen our understanding of how training data gaps and model architectures interact to produce systematic biases in individual contexts.

In conclusion, the foreign bias we document reveals how AI can produce unexpected behavioral patterns that systematically distort financial predictions. Rather than serving as neutral arbiters, AI appears to be introducing novel forms of bias that require careful study and active mitigation. As financial markets increasingly rely on AI-generated analysis, understanding and correcting these biases becomes essential for maintaining market efficiency and protecting investors.

References

Bender, E. M., T. Gebru, A. McMillan-Major, and S. Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? *FAccT 2021 - Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* 610–23.

Bertomeu, J., Y. Lin, Y. Liu, and Z. Ni. 2025. The impact of generative AI on information processing: Evidence from the ban of ChatGPT in Italy. *Journal of Accounting and Economics* 80:101782–.

Bini, P., L. W. Cong, X. Huang, and J. Lawrence J. 2025. Behavioral Economics of AI: LLM Biases and Corrections. *Working Paper* .

Blankespoor, E., E. DeHaan, and Q. Li. 2025. Generative AI in Financial Reporting. *Working Paper* 1–23.

Blodgett, S. L., S. Barocas, H. Daumé, and H. Wallach. 2020. Language (Technology) is power: A critical survey of "bias" in NLP. *Proceedings of the Annual Meeting of the Association for Computational Linguistics* 5454–76.

Bloomberg Intelligence. 2024. Generative AI 2024: Assessing Opportunities and Disruptions in an Evolving Trillion-Dollar Market. Working Paper.

Bradshaw, M. T., C. Ma, B. P. Yost, and Y. Zou. 2025. Generative AI Use by Capital Market Information Intermediaries : Evidence from Seeking Alpha. *Working Paper* .

Brennan, M. J., and A. Subrahmanyam. 1995. Investment analysis and price formation in securities markets. *Journal of Financial Economics* 38:361–81.

Bushee, B. J., J. E. Core, W. Guay, and S. J. Hamm. 2010. The role of the business press as an information intermediary. *Journal of Accounting Research* 48:1–19.

Cao, S., W. Jiang, J. Wang, and B. Yang. 2024. From Man vs. Machine to Man + Machine: The art and AI of stock analyses. *Journal of Financial Economics* 160:103910–.

Chen, S., T. C. Green, H. Gulen, and D. Zhou. 2024. What Does ChatGPT Make of Historical Stock Returns? Extrapolation and Miscalibration in LLM Stock Return Forecasts .

Cheng, Q., P. Lin, and Y. Zhao. 2025. Does Generative AI Facilitate Investor Trading? Evidence from ChatGPT Outages Does Generative AI Facilitate Investor Trading? Evidence from ChatGPT Outages. *Singapore Management University School of Accountancy Research Paper Series* 13:101821–.

Cook, T., S. Kazinnik, A. Hansen, and P. McAdam. 2023. Evaluating Local Language Models: An Application to Bank Earnings Calls. *The Federal Reserve Bank of Kansas City Research Working Papers* .

Cooper, I., and E. Kaplanis. 1994. Home Bias in Equity Portfolios, Inflation Hedging, and International Capital Market Equilibrium. *Review of Financial Studies* 7:45–60.

Coval, J. D., and T. J. Moskowitz. 1999. Home bias at home: Local equity preference in domestic portfolios. *Journal of Finance* 54:2045–73.

de Kok, T. 2025. ChatGPT for Textual Analysis? How to Use Generative LLMs in Accounting Research. *Management Science (Forthcoming)* .

Didisheim, A., M. Fraschini, and L. Somoza. 2025. Ai's predictable memory and its implications for finance. *Available at SSRN 5331177* .

Drake, M. S., N. M. Guest, and B. J. Twedt. 2014. The media and mispricing: The role of the business press in the pricing of accounting information. *Accounting Review* 89:1673–701.

Engelberg, J., A. Manela, W. Mullins, and L. Vulicevic. 2025. Entity Neutering. *Working Paper* 1–53.

French, K. R., and J. M. Poterba. 1991. Investor diversification and international equity markets.

Gao, Y., M. Zhang, and H. Yang. 2023. Looking good in the eyes of stakeholders: corporate giving and corporate acquisitions. *Journal of Business Ethics* 185:375–96.

Guest, N. M. 2021. The information role of the media in earnings news. *Journal of Accounting Research* 59:1021–76.

He, S., L. Lv, A. Manela, and J. Wu. 2025. Chronologically Consistent Large Language Models. *Working Paper* .

Heath, C., and A. Tversky. 1991. Preference and belief: Ambiguity and competence in choice under uncertainty. *Journal of Risk and Uncertainty* 4:5–28.

Huberman, G. 2001. Familiarity breeds investment. *Review of Financial Studies* 14:659–80.

Jha, M., J. Qian, M. Weber, and B. Yang. 2024. ChatGPT and Corporate Policies. *SSRN Electronic Journal* .

Kim, A., M. Muhn, and V. Nikolaev. 2024. Financial Statement Analysis with Large Language Models .

Kim, A. G., and V. V. Nikolaev. 2025. Context-Based Interpretation of Financial Information. *Journal of Accounting Research (Forthcoming)* .

Law, K. K., and M. Shen. 2025. How Does Artificial Intelligence Shape Audit Firms? *Management Science* 71:3641–66.

Lee, H., J. Seo, S. Park, J. Lee, W. Ahn, C. Choi, A. Lopez-Lira, and Y. Lee. 2025. *Your AI, Not Your View: The Bias of LLMs in Investment Analysis*, vol. 1. arXiv.

Levy, B. 2025. Caution Ahead: Numerical Reasoning and Look-ahead Bias in AI Models. *Working Paper* .

Lewis, K. K. 1999. Trying to Explain Home Bias in Equities and Consumption. *Journal of Economic Literature* 37:571–608.

Li, J., M. Li, X. Wang, and J. B. Thatcher. 2021. Strategic directions for ai: The role of cios and boards of directors. *MIS quarterly* .

Lock, B. 2024. The impact of news media coverage on voluntary disclosure. *Contemporary Accounting Research* 41:2354–83.

Lopez-Lira, A., and Y. Tang. 2023. Can ChatGPT Forecast Stock Price Movements? Return Predictability and Large Language Models. *SSRN Electronic Journal* .

Portes, R., and H. Rey. 2005. The determinants of cross-border equity flows. *Journal of International Economics* 65:269–96.

Qiao, F. 2023. Do analysts disseminate anomaly information in china? *Journal of Banking and Finance, Forthcoming* .

Sarkar, S., and K. Vafa. 2024. Lookahead Bias in Pretrained Language Models. *SSRN Electronic Journal* .

Shaffer, M., and C. C. Y. Wang. 2024. Scaling Core Earnings Measurement with Large Language Models. *Working Paper* .

Shumailov, I., Z. Shumaylov, Y. Zhao, Y. Gal, N. Papernot, and R. Anderson. 2023. The Curse of Recursion: Training on Generated Data Makes Models Forget .

Tesar, L. L., and I. M. Werner. 1995. Home bias and high turnover. *Journal of International Money and Finance* 14:467–92.

Van Nieuwerburgh, S., and L. Veldkamp. 2009. Information immobility and the home bias puzzle. *Journal of Finance* 64:1187–215.

Appendix A Description of Variables

This table defines variables used in our analyses. Variable definitions are based on data obtained from ChatGPT 4.1, DeepSeek R1, RavenPack News Analytics, CFND, Compustat, CSMAR, and Google, as described in Section 2.

Variable	Description	Source
	Key Variables of Interest	
AI_PREDICTION	Stock price prediction of gen AI for a Chinese firm for Dec 31, 2024.	ChatGPT 4.1, DeepSeek R1
AI_RECOMMEND	Indicator = 1 if gen AI gives a purchase recommendation for June 30–Dec 31, 2024, for a Chinese firm; 0 otherwise.	ChatGPT 4.1, DeepSeek R1
US_AI	Indicator = 1 if the gen AI is the US AI; 0 otherwise.	ChatGPT 4.1, DeepSeek R1
NEG_NEWS_GAP	the number of negative news from Chinese media minus the number of negative news from US media divided by their total number during the three-years period prior to the end of training period: i.e., 2022.06.30 to 2024.06.30. It is set to be zero if the value is missing.	RavenPack News Analytics, CFND
NET_NEG_NEWS_GAP	the number of negative news minus the number of positive news from Chinese media minus the number of negative news minus the number of positive news from US media, divided by their total number, during the three-years period prior to the end of training period: i.e., 2022.06.30 to 2024.06.30.	RavenPack News Analytics, CFND
	Firm-Level Control Variables	
FIRMSIZE	Log of total assets in year $t-1$.	Compustat
LEV	Year-end ratio of total liabilities to total assets in year $t-1$.	Compustat
ROA	Income before extraordinary items / total assets in year $t-1$.	Compustat
LOSS	Indicator = 1 if net income is negative in year $t-1$; 0 otherwise.	Compustat
BM	Book-to-market ratio in year $t-1$.	Compustat
PRE_RETURN	Annual stock returns in year $t-1$.	Compustat
RETURN_VOLATILITY	Std. dev. of daily stock returns in year $t-1$.	Compustat
IO	Percentage of institutional ownership in year $t-1$.	CSMAR

Variable	Description	Computation
Variables in Other Tests		
AI_PREDICTION_FEEDED	AI prediction for Dec 31, 2024, after feeding Chinese media news if available.	ChatGPT 4.1, DeepSeek R1, CFND
AI_RECOMMEND_FEEDED	Indicator = 1 if AI recommends purchase for June 30–Dec 31, 2024, after feeding Chinese media news; 0 otherwise.	ChatGPT 4.1, DeepSeek R1, CFND
FOREIGN_IO	Percentage of foreign institutional ownership in year $t-1$.	CSMAR
AI_PREDICTION_ERROR	|AI predicted price – actual price at Dec 31, 2024|.	ChatGPT 4.1, DeepSeek R1, Compustat
AI_RECOMMEND_ERROR	|AI recommendation value – actual recommendation|, where actual recommendation = price change from June 30 to Dec 31, 2024.	ChatGPT 4.1, DeepSeek R1, Compustat
AI_PREDICTION_US	AI prediction for US firm for Dec 31, 2024.	ChatGPT 4.1, DeepSeek R1
AI_RECOMMEND_US	Indicator = 1 if AI recommends purchase for June 30–Dec 31, 2024, for a US firm; 0 otherwise.	ChatGPT 4.1, DeepSeek R1
AI_PREDICTION_ALT1	AI prediction for Chinese firm for June 30, 2025.	ChatGPT 4.1, DeepSeek R1
AI_RECOMMEND_ALT1	Indicator = 1 if AI recommends purchase for June 30, 2024–June 30, 2025, for a Chinese firm; 0 otherwise.	ChatGPT 4.1, DeepSeek R1
AI_PREDICTION_ALT2	AI prediction for Chinese firm for June 30, 2025, *using web search*.	ChatGPT 4.1, DeepSeek R1
AI_RECOMMEND_ALT2	Indicator = 1 if AI recommends purchase for June 30, 2024–June 30, 2025, for a Chinese firm, *using web search*; 0 otherwise.	ChatGPT 4.1, DeepSeek R1
NEG_NEWS_GAP_2019.22	the number of negative news from Chinese media minus the number of negative news from US media divided by their total number during the three-years period prior to the end of training period: i.e., 2019.11.30 to 2022.11.30. It is set to be zero if the value is missing.	RavenPack News Analytics, CFND
NET_NEG_NEWS_GAP_2019.22	the number of negative news minus the number of positive news from Chinese media minus the number of negative news minus the number of positive news from US media, divided by their total number, during the three-years period prior to the end of training period: i.e., 2019.11.30 to 2022.11.30.	RavenPack News Analytics, CFND
NEG_NEWS_GAP_2022.25	the number of negative news from Chinese media minus the number of negative news from US media divided by their total number during the three-years period prior to the end of training period: i.e., 2022.01.20 to 2025.01.20. It is set to be zero if the value is missing.	RavenPack News Analytics, CFND
NET_NEG_NEWS_GAP_2022.25	the number of negative news minus the number of positive news from Chinese media minus the number of negative news minus the number of positive news from US media, divided by their total number, during the three-years period prior to the end of training period: i.e., 2022.01.20 to 2025.01.20.	RavenPack News Analytics, CFND

Table 1. Summary Statistics

This table reports descriptive statistics for the main variables in our sample of Chinese firms used in the baseline analysis. Panel A provides the distribution of the primary dependent and independent variables across 9,956 firm–AI observations, including AI price predictions, buy recommendations, and key firm characteristics. Panel B presents the average news gap measures by two-digit GICS industry, showing substantial cross-industry variation in the extent of U.S.–China media coverage asymmetries. All variables are defined in Appendix A and all continuous variables are winsorized at the 1st and 99th percentiles.

Panel A: Distribution of Main Variables

Variable	*N*	*Mean*	*SD*	*p10*	*p25*	*p50*	*p75*	*p90*
AI_PREDICTION	9,956	20.76	16.07	5.33	8.32	15.33	28.46	45.33
AI_RECOMMEND	9,956	0.98	0.13	1.00	1.00	1.00	1.00	1.00
US_AI	9,956	0.50	0.50	0.00	0.00	0.50	1.00	1.00
NEG_NEWS_GAP	9,956	0.58	0.40	-0.03	0.33	0.50	1.00	1.00
NET_NEG_NEWS_GAP	9,956	0.05	0.18	-0.10	-0.07	0.00	0.09	0.26
FIRM_SIZE	9,956	8.60	1.40	7.13	7.62	8.29	9.30	10.50
ROA	9,956	0.02	0.06	-0.05	0.00	0.03	0.05	0.09
LEVERAGE	9,956	0.42	0.22	0.13	0.24	0.40	0.57	0.71
LOSS	9,956	0.21	0.41	0.00	0.00	0.00	0.00	1.00
BM	9,956	0.57	0.46	0.18	0.29	0.45	0.70	1.02
PRE_RETURN	9,956	0.03	0.32	-0.32	-0.17	0.00	0.17	0.39
RETURN_VOLATILITY	9,956	2.36	0.88	1.36	1.71	2.19	2.88	3.59
IO	9,956	41.11	25.12	7.28	19.63	40.43	61.78	75.79

Panel B: Average News Gap by Fama-French 12-industry classification

Industry	*NEG_NEWS_GAP*	*NET_NEG_NEWS_GAP*
Consumer Nondurables	0.6571	0.0430
Consumer Durables	0.5844	0.0477
Manufacturing	0.5902	0.0307
Energy	0.5417	0.0402
Chemicals	0.6115	0.0238
Business Equipment	0.5871	0.0650
Telecommunications	0.6084	0.0388
Utilities	0.6326	0.0405
Wholesale and Retail	0.6022	0.0666
Healthcare	0.6281	0.0465
Financials	0.2574	0.2245
Others	0.5799	0.0527

Table 2. ChatGPT vs. DeepSeek: Price Predictions and Stock Recommendations

This table presents baseline regressions comparing six-month-ahead stock price predictions (Column 1) and buy recommendations (Column 2) for Chinese firms generated by ChatGPT (US-based) and DeepSeek (China-based). The key independent variable, *US AI*, equals one for ChatGPT predictions and zero for DeepSeek. Control variables include firm size, leverage, profitability, book-to-market ratio, prior returns, return volatility, loss indicator, and institutional ownership, with industry fixed effects. The coefficients on *US AI* capture the magnitude of the documented "foreign bias." All variables are defined in Appendix A and all continuous variables are winsorized at the 1st and 99th percentiles. Standard errors, clustered at the two-digit GICS industry level, are reported in parentheses. Significance levels are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

	(1)	(2)
	AI_PREDICTION	*AI_RECOMMEND*
US_AI	2.4153***	0.0135***
	(0.722)	(0.004)
FIRM_SIZE	0.7579	-0.0025
	(0.788)	(0.002)
ROA	33.7108***	-0.0187
	(7.558)	(0.054)
LEVERAGE	-15.9897***	-0.0235***
	(3.649)	(0.007)
LOSS	0.6487	-0.0114
	(1.013)	(0.012)
BM	-4.1948***	-0.0349***
	(1.216)	(0.009)
PRE_RETURN	-8.6622***	0.0008
	(0.951)	(0.008)
RETURN_VOLATILITY	4.9367***	-0.0009
	(0.778)	(0.002)
IO	0.0702*	0.0001
	(0.033)	(0.000)
Industry FE	Yes	Yes
Observations	9,956	9,956
Adj R^2	0.3071	0.0325

Table 3. Role of Media Coverage in AI Prediction: Gap in News Coverage

This table examines whether cross-border media coverage asymmetries explain the foreign bias in AI-generated predictions. We regress predicted prices (Columns 1–2) and buy recommendations (Columns 3–4) on the interaction between *US AI* and news gap measures: *NEG NEWS GAP* and *NET NEG NEWS GAP*. The news gap variables measure the relative prevalence of U.S. versus Chinese negative news coverage for each firm. Negative interaction coefficients indicate that greater U.S. coverage—particularly of negative news—reduces ChatGPT's relative optimism. All continuous variables are winsorized at the 1st and 99th percentiles. All variables are defined in Appendix A and all continuous variables are winsorized at the 1st and 99th percentiles. Standard errors, clustered at the Fama-French 12-industry level, are reported in parentheses. Significance levels are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

	(1)	(2)	(3)	(4)
	AI_PREDICTION		*AI_RECOMMEND*	
US_AI×*NEG_NEWS_GAP*	3.5373***		0.0439*	
	(0.507)		(0.023)	
US_AI×*NET_NEG_NEWS_GAP*		30.2682***		0.1420***
		(7.466)		(0.023)
NEG_NEWS_GAP	-4.3225***		-0.0311	
	(0.828)		(0.023)	
NET_NEG_NEWS_GAP		0.8512		-0.1265***
		(3.140)		(0.030)
US_AI	1.1700	-0.7362	-0.0092	0.0113
	(0.757)	(0.511)	(0.008)	(0.006)
FIRM_SIZE	0.6437	0.6022	-0.0032	-0.0022
	(0.770)	(0.740)	(0.002)	(0.002)
ROA	34.6185***	34.3042***	-0.0172	-0.0305
	(7.389)	(6.347)	(0.051)	(0.053)
LEVERAGE	-15.5983***	-14.9733***	-0.0221***	-0.0232***
	(3.634)	(3.048)	(0.007)	(0.006)
LOSS	0.7682	0.6404	-0.0113	-0.0117
	(0.975)	(0.946)	(0.011)	(0.012)
BM	-4.0289***	-4.0166***	-0.0336***	-0.0343***
	(1.213)	(1.057)	(0.009)	(0.009)
PRE_RETURN	-8.7422***	-8.1362***	0.0007	0.0016
	(0.919)	(0.846)	(0.007)	(0.008)
RETURN_VOLATILITY	4.9618***	4.6252***	-0.0008	-0.0012
	(0.770)	(0.725)	(0.002)	(0.002)
IO	0.0731**	0.0720**	0.0001	0.0001
	(0.033)	(0.028)	(0.000)	(0.000)
Industry FE	Yes	Yes	Yes	Yes
Observations	9,956	9,956	9,956	9,956
Adj R^2	0.3102	0.3590	0.0354	0.0455

Table 4. Role of Media Coverage: When Firm Has No News

This table investigates whether foreign bias persists when both U.S. and Chinese media coverage is absent, and whether placebo firms show similar patterns. Panel A restricts the sample to firms with no relevant U.S. or Chinese news in the past three years. Panel B reports results for a placebo set of synthetic Chinese firm names. All continuous variables are winsorized at the 1st and 99th percentiles. All variables are defined in Appendix A and all continuous variables are winsorized at the 1st and 99th percentiles. Standard errors, clustered at the Fama-French 12-industry level, are reported in parentheses. Significance levels are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

Panel A: Actual Firms with No U.S. and No Chinese News

	(1)	(2)
	AI_PREDICTION	*AI_RECOMMEND*
US_AI	1.6741	0.0000
	(1.455)	(0.016)
FIRM_SIZE	0.9747	-0.0069
	(1.102)	(0.013)
ROA	39.8488	-0.6682
	(28.226)	(0.490)
LEVERAGE	-21.6819***	-0.0316
	(6.484)	(0.051)
LOSS	6.0285	-0.1562
	(5.484)	(0.108)
BM	-0.6179	-0.0575*
	(1.052)	(0.026)
PRE_RETURN	-6.2851***	0.0068
	(1.751)	(0.055)
RETURN_VOLATILITY	4.4628**	0.0028
	(1.762)	(0.023)
IO	0.0789*	0.0009
	(0.036)	(0.001)
Industry FE	Yes	Yes
Observations	382	382
Adj R^2	0.3944	0.0772

Table 2 [Continued]

Panel B: Placebo (Synthetic Chinese Firm, No News) Sample

	(1) *AI_PREDICTION*	(2) *AI_RECOMMEND*
US_AI	-0.2260	0.0100
	(0.275)	(0.029)
FIRM_SIZE	-0.5009	-0.0379
	(0.801)	(0.023)
ROA	-47.7204*	-0.2994
	(22.845)	(0.423)
LEVERAGE	0.3245	0.0444
	(6.373)	(0.108)
LOSS	-9.1091**	-0.1418*
	(3.396)	(0.066)
BM	0.1858	0.0306
	(1.098)	(0.054)
PRE_RETURN	-4.4750	-0.0586
	(3.371)	(0.085)
RETURN_VOLATILITY	2.5929**	0.0001
	(0.863)	(0.037)
IO	0.0567	0.0016
	(0.036)	(0.001)
Industry FE	Yes	Yes
Observations	200	200
Adj R^2	0.1773	0.0909

Table 5. Closing the News Gap

This table reports results from the information intervention in which ChatGPT was provided with up to five recent Chinese news articles about each firm prior to making predictions. The intervention eliminates the foreign bias: the *US AI* coefficient becomes statistically indistinguishable from zero for both price predictions (Column 1) and buy recommendations (Column 2). This finding supports the mechanism that the bias is driven by asymmetric information processing. All variables are defined in Appendix A and all continuous variables are winsorized at the 1st and 99th percentiles. Standard errors, clustered at the Fama-French 12-industry level, are reported in parentheses. Significance levels are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

	(1)	(2)
	AI_PREDICTION_FEEDED	*AI_RECOMMEND_FEEDED*
US_AI	-0.7980	-0.0022
	(0.770)	(0.003)
FIRM_SIZE	0.6187	-0.0047
	(0.808)	(0.003)
ROA	33.5798***	-0.0372
	(8.068)	(0.069)
LEVERAGE	-16.5883***	-0.0251***
	(3.731)	(0.005)
LOSS	0.4735	-0.0132
	(1.037)	(0.015)
BM	-4.0468***	-0.0601***
	(1.211)	(0.014)
PRE_RETURN	-9.0368***	-0.0039
	(0.986)	(0.010)
RETURN_VOLATILITY	5.1877***	-0.0014
	(0.818)	(0.002)
IO	0.0752*	0.0002*
	(0.035)	(0.000)
Industry FE	Yes	Yes
Observations	9,956	9,956
Adj R^2	0.3060	0.0492

Table 6. Boundary Condition - US firms (placebo test)

This table repeats the baseline analysis for a sample of 11,688 U.S. firm-year observations. The dependent variables are price predictions (Column 1) and buy recommendations (Column 2) for U.S. firms, with *US AI* indicating ChatGPT predictions. The coefficients on *US AI* are small and statistically insignificant, confirming that the documented foreign bias is specific to cross-border predictions about Chinese firms. All continuous variables are winsorized at the 1st and 99th percentiles. All variables are defined in Appendix A. Standard errors, clustered at the Fama-French 12-industry level, are reported in parentheses. Significance levels are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

	(1)	(2)
	AI_PREDICTION_US	*AI_RECOMMEND_US*
US_AI	-0.1539	-0.0472
	(0.361)	(0.027)
FIRM_SIZE	10.8241***	0.0191***
	(2.254)	(0.004)
ROA	-3.2749***	0.0183***
	(0.674)	(0.004)
LEVERAGE	-0.0487	-0.0030
	(0.183)	(0.003)
LOSS	-31.9823***	-0.0500**
	(5.347)	(0.022)
BM	-4.6968***	-0.0017
	(0.931)	(0.004)
PRE_RETURN	0.7273***	0.0061*
	(0.184)	(0.003)
RETURN_VOLATILITY	-0.0857	-0.0099**
	(0.087)	(0.004)
IO	12.9281	-0.0110
	(11.715)	(0.013)
Industry FE	Yes	Yes
Observations	11,688	11,688
Adj R^2	0.2740	0.1759

Table 7. Prediction Errors

This table compares the accuracy of ChatGPT and DeepSeek predictions against realized outcomes. Column 1 reports absolute price prediction errors, and Column 2 reports recommendation errors, defined as the absolute deviation from an ex-post optimal buy/hold/sell classification based on realized returns. Positive coefficients on *US AI* indicate that ChatGPT is less accurate than DeepSeek, consistent with its optimism reflecting bias rather than superior information. All variables are defined in Appendix A and all continuous variables are winsorized at the 1st and 99th percentiles. Standard errors, clustered at the Fama-French 12-industry level, are reported in parentheses. Significance levels are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

	(1)	(2)
	AI_PREDICTION_ERRORS	*AI_RECOMMEND_ERRORS*
US_AI	1.0863***	0.0074*
	(0.318)	(0.004)
FIRM_SIZE	0.3676	0.0121
	(0.454)	(0.013)
ROA	24.6574***	0.9947***
	(5.694)	(0.155)
LEVERAGE	-6.2347***	0.0307
	(1.926)	(0.029)
LOSS	1.8352**	0.0532***
	(0.657)	(0.015)
BM	-2.8381**	0.0875**
	(0.946)	(0.030)
PRE_RETURN	-1.3305	0.0206
	(0.815)	(0.035)
RETURN_VOLATILITY	2.6255***	0.0431**
	(0.566)	(0.014)
IO	0.0462*	0.0003
	(0.022)	(0.000)
Industry FE	Yes	Yes
Observations	9,956	9,956
Adj R^2	0.1487	0.0565

Table 8. Robustness Tests

This table presents three sets of robustness analyses. Panel A extends the prediction horizon to 12 months, showing the foreign bias persists for longer horizons. Panel B enables web search for both models before prediction, which reduces but does not eliminate the bias. Panel C repeats the baseline prompts in Mandarin, showing the bias remains, albeit somewhat attenuated. Together, these tests confirm that the foreign bias is robust to alternative prediction windows, real-time information access, and prompt language. All continuous variables are winsorized at the 1st and 99th percentiles. All variables are defined in Appendix A and all continuous variables are winsorized at the 1st and 99th percentiles. Standard errors, clustered at the Fama-French 12-industry level, are reported in parentheses. Significance levels are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

Panel A: Alternative Prediction Horizon

	(1)	(2)
	AI_PREDICTION_ALT1	*AI_RECOMMEND_ALT1*
US_AI	1.8490**	0.0207***
	(0.818)	(0.006)
FIRM_SIZE	4.2758	-0.0059
	(2.668)	(0.005)
ROA	71.6423**	-0.0357
	(24.110)	(0.094)
LEVERAGE	-26.2111***	-0.0284**
	(8.427)	(0.012)
LOSS	6.1597	-0.0283
	(3.553)	(0.020)
BM	-10.1748*	-0.0453**
	(4.901)	(0.016)
PRE_RETURN	-9.3972***	-0.0084
	(1.374)	(0.011)
RETURN_VOLATILITY	4.8821***	0.0041
	(1.037)	(0.003)
IO	0.0624	0.0000
	(0.039)	(0.000)
Industry FE	Yes	Yes
Observations	9,956	9,956
Adj R^2	0.1037	0.0479

Table 8 [Continued]

Panel B: AI Prediction Allowing Online Search

	(1)	(2)
	AI_PREDICTION_ALT2	*AI_RECOMMEND_ALT2*
US_AI	0.5338**	0.0426**
	(0.228)	(0.016)
FIRM_SIZE	0.9237***	0.0072
	(0.151)	(0.013)
ROA	2.0662	-1.2009***
	(2.513)	(0.212)
LEVERAGE	-0.4815	0.3108***
	(0.455)	(0.067)
LOSS	0.9130***	0.0068
	(0.249)	(0.021)
BM	-1.3104***	0.0653**
	(0.321)	(0.024)
PRE_RETURN	0.2101	0.1196***
	(0.357)	(0.029)
RETURN_VOLATILITY	-0.5043**	-0.1436***
	(0.164)	(0.021)
IO	-0.0077*	-0.0015*
	(0.004)	(0.001)
Industry FE	Yes	Yes
Observations	9,956	9,956
Adj R^2	0.0186	0.2017

Table 8 [Continued]

Panel C: AI Prediction with Prompt in Mandarin

	(1)	(2)
	AI_PREDICTION_ALT3	*AI_RECOMMEND_ALT3*
US_AI	1.5173**	0.0135*
	(0.644)	(0.007)
FIRM_SIZE	3.8150	0.0157***
	(2.549)	(0.004)
ROA	60.9911**	0.0607
	(23.239)	(0.077)
LEVERAGE	-22.6838**	-0.0342*
	(7.758)	(0.017)
LOSS	5.5226	-0.0068
	(3.226)	(0.019)
BM	-7.3221	0.0498***
	(4.620)	(0.016)
PRE_RETURN	-7.4429***	0.0116
	(1.305)	(0.008)
RETURN_VOLATILITY	4.0061***	0.0019
	(0.928)	(0.004)
IO	0.0553	-0.0003**
	(0.039)	(0.000)
Industry FE	Yes	Yes
Observations	9,956	9,956
Adj R^2	0.0763	0.0175

Table 9. Market reactions

This table reports exploratory event study regressions examining cumulative abnormal returns (CARs) around the public release of ChatGPT 3.5 on November 30, 2022 (Panel A) and DeepSeek R1 on January 20, 2025 (Panel B). All variables are defined in Appendix A and all continuous variables are winsorized at the 1st and 99th percentiles. Standard errors, clustered at the Fama-French 12-industry level, are reported in parentheses. Significance levels are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

Panel A: ChatGPT 3.5 Announcement

	(1)	(2)	(3)	(4)	(5)	(6)
	CAR(-1,1)	*CAR(-3,3)*	*CAR(-1,1)*	*CAR(-3,3)*	*CAR(-1,1)*	*CAR(-3,3)*
IO_FOREIGN	0.0031**	0.0130***	-0.0023	0.0043	0.0047**	0.0135***
	(0.001)	(0.005)	(0.002)	(0.003)	(0.002)	(0.004)
IO_FOREIGN×*NEG_NEWS_GAP_2019_22*			0.0145***	0.0236**		
			(0.006)	(0.012)		
IO_FOREIGN×*NET_NEG_NEWS_GAP_2019_22*					0.0107**	0.0215*
					(0.005)	(0.012)
NEG_NEWS_GAP_2019_22			0.0083***	0.0113***		
			(0.002)	(0.004)		
NET_NEG_NEWS_GAP_2019_22					-0.0219***	-0.0091**
					(0.002)	(0.004)
FIRM_SIZE	0.0024***	0.0041***	0.0028***	0.0046***	0.0015**	0.0037***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
ROA	-0.0052	-0.0731***	-0.0120	-0.0829***	-0.0013	-0.0713***
	(0.015)	(0.027)	(0.015)	(0.027)	(0.015)	(0.027)
LEVERAGE	-0.0110***	-0.0029	-0.0122***	-0.0045	-0.0091**	-0.0029
	(0.004)	(0.007)	(0.004)	(0.007)	(0.004)	(0.007)
LOSS	-0.0010	-0.0036	-0.0016	-0.0045	-0.0002	-0.0034
	(0.002)	(0.004)	(0.002)	(0.004)	(0.002)	(0.004)
BM	-0.0094***	-0.0111***	-0.0095***	-0.0112***	-0.0085***	-0.0111***
	(0.002)	(0.003)	(0.002)	(0.003)	(0.002)	(0.003)
PRE_RETURN	-0.0015	0.0207***	-0.0012	0.0211***	-0.0013	0.0202***
	(0.002)	(0.003)	(0.002)	(0.003)	(0.002)	(0.003)
RETURN_VOLATILITY	0.0008	-0.0059***	0.0007	-0.0060***	0.0006	-0.0057***
	(0.001)	(0.002)	(0.001)	(0.002)	(0.001)	(0.001)
IO	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4,702	4,702	4,702	4,702	4,702	4,702
Adj R^2	0.0396	0.0704	0.0573	0.0822	0.0793	0.0815

Table 9 [Continued]

Panel B: DeepSeek R1 Announcement

	(1)	(2)	(3)	(4)	(5)	(6)
	CAR(-1,1)	*CAR(-3,3)*	*CAR(-1,1)*	*CAR(-3,3)*	*CAR(-1,1)*	*CAR(-3,3)*
IO_FOREIGN	-0.0013***	-0.0020**	0.0008	0.0017	-0.0027**	-0.0039**
	(0.000)	(0.001)	(0.001)	(0.002)	(0.001)	(0.002)
IO_FOREIGN×*NEG_NEWS_GAP_2022_25*			-0.0032**	-0.0056**		
			(0.002)	(0.003)		
IO_FOREIGN×*NET_NEG_NEWS_GAP_2022_25*					-0.0041**	-0.0056*
					(0.002)	(0.003)
NEG_NEWS_GAP_2022_25			-0.0008	-0.0025	-0.0012	-0.0032*
			(0.001)	(0.002)	(0.001)	(0.002)
FIRM_SIZE	-0.0020***	-0.0079***	-0.0020***	-0.0080***	-0.0021***	-0.0081***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
ROA	0.0377**	0.0170	0.0384**	0.0190	0.0379**	0.0183
	(0.015)	(0.022)	(0.015)	(0.022)	(0.015)	(0.022)
LEVERAGE	-0.0074**	0.0090*	-0.0073**	0.0093*	-0.0070*	0.0099*
	(0.004)	(0.005)	(0.004)	(0.005)	(0.004)	(0.005)
LOSS	-0.0069***	-0.0040	-0.0068***	-0.0038	-0.0067***	-0.0037
	(0.002)	(0.003)	(0.002)	(0.003)	(0.002)	(0.003)
BM	-0.0048***	0.0135***	-0.0049***	0.0135***	-0.0046***	0.0138***
	(0.001)	(0.002)	(0.001)	(0.002)	(0.001)	(0.002)
PRE_RETURN	0.0014	0.0059**	0.0013	0.0056**	0.0014	0.0058**
	(0.002)	(0.003)	(0.002)	(0.003)	(0.002)	(0.003)
RETURN_VOLATILITY	0.0014**	0.0009	0.0015**	0.0010	0.0014**	0.0009
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
IO	-0.0000	-0.0001*	-0.0000	-0.0001	-0.0000	-0.0000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4,978	4,978	4,978	4,978	4,978	4,978
Adj R^2	0.0650	0.0353	0.0656	0.0369	0.0671	0.0377