

Risk and Return under Open-Source LLM Shocks: A CAPM Analysis of U.S. AI-Related Sector ETFs

Litao Zhou

*Business College, Wuxi Taihu University, Wuxi, China
222042215@wxu.edu.cn*

Abstract. The rapid iteration of open-source large language models is reshaping the competitive landscape of the AI industry, yet systematic empirical evidence on their specific effects on capital markets remains limited. Using the CAPM as the analytical framework and combining it with the event-study methodology, this paper examines the impact of two landmark open-source releases—DeepSeek-R1 and Kimi K3—on the returns and risk of five U.S. AI-related sector ETFs. Three findings emerge: First, both shocks were highly concentrated in the computing-power layer, while the application layer remained relatively resilient, presenting a short-term pattern of repricing within the AI industry chain. Second, the Kimi K3 shock produced shallower losses and a faster recovery, which may indicate that the market had developed a certain degree of immunity to such events. Third, following the DeepSeek shock, the risk structure of the sectors underwent changes that persisted throughout the sample period: the beta of SOXX increased, whereas the betas of software, communication services, and utilities declined markedly, with AI-related systematic risk exhibiting a tendency to concentrate in the core semiconductor segment.

Keywords: CAPM, open-source LLMs, sector ETFs

1. Introduction

In the current era, AI is developing at a remarkable pace, profoundly transforming social life and reshaping a wide range of industries. Among its various branches, the iteration of LLMs is particularly striking: their capability boundaries continue to expand and now cover conversational interaction, content generation, programming assistance, and other domain-specific tasks. To date, the LLM industry has been highly favored by capital markets, with leading firms such as Anthropic and OpenAI initiating their listing processes and valuations approaching one trillion U.S. dollars, painting a seemingly prosperous picture of the industry's prospects. However, it is worth noting that LLM companies generally exhibit limited profitability, and the industry is still exploring sustainable business models. Against this backdrop, continuous capital investment has become critical to the survival and development of LLM firms. The rise of open-source large language models (LLMs) has reshaped this competitive landscape in tangible ways: DeepSeek-R1's 2025 launch stunned the industry with its unprecedentedly low training costs, while Kimi K3's 2026 release reignited fervent discussions across the AI community. When high-performance models are made freely available through open source, investors' willingness to invest may be affected, with repercussions potentially

extending to AI-related industries such as semiconductors. However, existing research on the market impact of large models has focused mainly on the firm level, while sector ETFs—the primary investment instruments actually held by investors—have received insufficient attention, and the impact of open-source releases on industry-level returns and risk structure remains to be examined. Therefore, using daily data on U.S. AI-related sector ETFs from 2021 to 2026, this paper applies the Capital Asset Pricing Model to investigate market reactions surrounding the two major releases of DeepSeek-R1 and Kimi K3, evaluating the market impact of open-source LLMs along two dimensions: short-term return shocks and long-term risk structure. The significance of this study lies in providing risk-identification and allocation references for investors with AI-sector exposure, as well as market-level empirical evidence for policymakers concerned with artificial intelligence industrial policy.

2. Literature review

Research on the Capital Asset Pricing Model (CAPM), both in China and abroad, is multifaceted and can be broadly divided into two strands. One strand empirically tests the applicability of the CAPM itself or proposes optimizations and improvements to the model. Ayub et al. constructed a six-factor DCAPM that replaces the traditional beta with a downside beta and incorporates a momentum factor, empirical evidence from the Pakistani stock market shows that it outperforms the traditional CAPM in both goodness of fit and the risk–return relationship, with particularly pronounced advantages during periods of high volatility [1]. Brusov et al. analytically embedded the CAPM and Fama–French models into the BFO and MM capital structure theories, constructing the CAPM-BFO and CAPM-MM valuation frameworks and demonstrating the fallacy of the Hamada levered CAPM formula [2]. Habis extended the theoretical foundation of the CAPM from two periods to three, proving that the beta pricing formula remains valid in a three-period general equilibrium model with complete markets and that equilibrium consumption allocations are Pareto efficient—a result indicating that the CAPM can be applied to long-term asset pricing [3]. Pooja et al. combined the CAPM with K-means and hierarchical clustering and found that the two approaches can stably group large-cap stocks by their risk–return characteristics, thereby supporting differentiated portfolio recommendations [4]. Ai and Zhang, using monthly data of 30 companies from 2002 to 2017, examined the applicability of the Sharpe–Lintner CAPM through full-sample, subsample, and portfolio-level tests, the results show that the CAPM generally holds at the individual-stock level (with insignificant alpha) but is significantly rejected at the portfolio level, indicating that beta cannot capture the pricing risk in portfolios driven by omitted factors such as size and value [5].

The other strand applies the CAPM as a tool for empirical testing and investment practice across different industries and market contexts. Hundal et al., based on 90 stocks listed on the Helsinki Stock Exchange from 2012 to 2016, showed that the risk–return relationship in the Finnish stock market is broadly consistent with the CAPM, with most stocks exhibiting positive Jensen's alphas, however, firm-specific risk accounts for approximately 80% of total risk, and outlier stocks can significantly distort the beta–return relationship [6]. Wijaya and Ferrari empirically tested 40 banking stocks on the Indonesia Stock Exchange using the CAPM and identified 31 efficient stocks and 9 inefficient ones by comparing actual returns with expected returns, thereby providing a basis for investment decisions in banking stocks [7]. Xiao, using daily data of 49 U.S. industry portfolios, compared the validity of the CAPM before and after the outbreak of COVID-19, the results show that, except for the gold industry, all industries support the CAPM both before and after the pandemic, with improved model fit and generally rising beta coefficients after the pandemic,

indicating a market characterized by the coexistence of high risk and high return [8]. Cao applied the CAPM to examine the relationship between risk and expected return for four energy commodity ETFs (USO, BNO, UNG, and UGA) from 2017 to 2021, finding that assets with positive betas all delivered positive expected returns, while UNG, with a negative beta, generated negative returns—a pattern consistent with a positive risk–return relationship, the study also noted that with annual data, the CAPM's pricing performance for natural gas ETFs deviates markedly, suggesting that data frequency and instrument heterogeneity require careful attention when applying the model to energy commodities [9].

In summary, existing research has, on the one hand, continuously tested and refined the CAPM itself and, on the other hand, applied it extensively to banking, energy, cross-national markets, and shock scenarios such as the pandemic. However, existing application-oriented studies have focused mainly on traditional industries and macro-level shocks, research targeting new technological shocks—particularly open-source LLM releases—remains limited, and most studies stay at the level of return performance around the event, paying insufficient attention to whether such shocks alter the industry's long-term risk structure (beta). This paper addresses this gap by applying the CAPM to U.S. AI-related sector ETFs and examining the short-term return shocks and long-term risk-structure changes associated with two open-source LLM release events.

3. Methodology

3.1. Population and sample

This study is quantitative research based on secondary data. The population of this study consists of five sector ETFs listed in the United States: SOXX (semiconductors), IGV (software), XLK (technology), XLC (communication services), and XLU (utilities). The sampling logic is as follows: the first four ETFs correspond respectively to the computing-power layer, the application layer, the broad technology layer, and the platform layer of the AI industry chain, while XLU serves as a defensive control sector with low relevance to AI. The sample period spans from August 2021 to August 2026, covering the two open-source LLM release events of DeepSeek-R1 and Kimi K3.

3.2. Data sources

The data used in this study are secondary data collected from the Tonghuashun iFinD financial data platform, including the daily adjusted closing prices of the five sector ETFs mentioned above. The market benchmark is proxied by the daily closing prices of SPY, and the risk-free rate is proxied by the daily yield of the U.S. 3-month Treasury, both likewise sourced from iFinD. The study period runs from August 2, 2021 to August 28, 2026, yielding 1,274 daily return observations. All data processing and estimation in this study were conducted in Python, using the pandas library for data handling and the statsmodels library for OLS regression.

3.3. Theoretical framework

This paper mainly applies the Capital Asset Pricing Model (CAPM) to investigate the impact of open-source LLM releases on capital-market sectors. According to Sharpe, the CAPM is given by [10]:

$$E(R_i) = R_f + \beta_i [E(R_m) - R_f] \quad (1)$$

where $E(R_i)$ denotes the expected return of security i , R_f is the risk-free rate, $E(R_m)$ is the expected return of the market portfolio, and β_i measures the systematic risk of security i . In this paper, R_f is proxied by the yield on U.S. 3-month Treasury securities, R_m by the daily return of SPY, and $i = 1, \dots, 5$ corresponds to the five sector ETFs. Empirically, this paper runs OLS regressions of each ETF's excess returns on daily frequency to estimate its α_i and β_i , the market shocks brought by the open-source releases are measured by abnormal returns, defined as the difference between actual returns and model-predicted values around the event dates. The changes in each sector's β before and after the DeepSeek event are further compared to assess the shock's impact on the long-term risk structure.

4. Results and discussion

4.1. Performance of the five sector etfs relative to the market

Table 1 summarizes the descriptive statistics of the five sector ETFs over the full sample period, covering six indicators: annualized return, annualized volatility, extreme daily returns (maximum and minimum), the correlation with SPY and cumulative return. Figure 1 displays the cumulative value of an initial \$100 investment in each sector over the five-year period. The results show that semiconductors represent the highest-return and highest-risk sector: SOXX achieved an annualized return of 31.83% with an annualized volatility of 38.28%—the latter being approximately 2.2 times that of the utilities sector (17.29%)—and a cumulative return of 245.32%, the highest among all sectors. XLK exhibits the next tier of high-risk, high-return characteristics (annualized return of 21.62% and volatility of 25.78%). By contrast, the cumulative returns of XLC and IGV were only 44.69% and 35.82%, respectively, even lower than the 49.42% of the defensive sector XLU. The correlations between each sector and SPY range from 0.44 (XLU) to 0.91 (XLK), indicating that XLU is suitable as a non-AI control sector. Figure 1 further shows that both open-source model releases were accompanied by notable drawdowns in cumulative value: following the release of DeepSeek-R1, SOXX experienced a sharp decline, then gradually recovered over the subsequent months and reached new highs, the drawdown after the Kimi K3 release was smaller in magnitude and shorter in duration. Taken together, these facts preliminarily indicate that the impact of open-source shocks is concentrated in high-beta sectors.

Table 1. Descriptive statistics of returns for the five sector ETFs

ETF	Return (%)	Volatility (%)	Min (%)	Max (%)	Corr. with SPY	Cumulative return (%)
SOXX	31.83	38.28	-10.44	18.57	0.80	245.32
IGV	10.09	28.43	-6.15	11.65	0.79	35.82
XLK	21.62	25.78	-6.82	13.43	0.91	152.26
XLC	9.48	20.83	-6.69	8.84	0.82	44.69
XLU	9.44	17.29	-5.56	4.71	0.44	49.42

Note: Return and volatility are annualized, Min and Max are the extreme daily returns over the sample period.




Figure 1. Five-year cumulative return of a \$100 investment in five U.S. sector ETFs

4.2. Full-sample CAPM regression results

Table 2 presents the full-sample CAPM regression results for the five sector ETFs, and Figure 2 plots them on the Security Market Line (SML). The betas of all five ETFs are highly significant at the 1% level, with SOXX exhibiting the highest beta (1.7943) and XLU the lowest (0.4409), the latter displaying clear defensive properties. The explanatory power of the single market factor for daily return fluctuations (R^2) ranges from 19.0% (XLU) to 83.5% (XLK). The AI-related technology sectors show generally higher betas and R^2 values than the defensive sector, indicating that their returns are more deeply linked to the market. In terms of positions relative to the SML, SOXX lies far above the line, IGV and XLC lie below it, and XLU sits close to the line. This indicates that the actual annualized excess returns of XLC and IGV are lower than the returns warranted by their risk levels—in other words, they bear market risk without receiving commensurate compensation. SOXX shows the opposite pattern, with its actual return exceeding the fair return by approximately 9.8%. Nevertheless, the absolute t-values of the alphas of all five ETFs are below 1, meaning that none of the alphas is statistically significant, which suggests that the CAPM prices these sector exposures reasonably well during normal periods.

Table 2. CAPM regression results

ETF	α (daily, %)	α (annualized, %)	$t(\alpha)$	β	$t(\beta)$	R^2
SOXX	0.0388	9.7880	0.96	1.7943	47.78	0.642
IGV	-0.0281	-7.0799	-0.91	1.3155	46.13	0.626
XLK	0.0152	3.8187	0.82	1.3778	80.14	0.835
XLC	-0.0179	-4.5106	-0.86	1.0027	51.67	0.677
XLU	0.0046	1.1671	0.17	0.4409	17.28	0.190




Figure 2. Security market line (SML) of the sector ETFs

4.3. Short-term abnormal returns around the two open-source releases

In the left panel (DeepSeek-R1), the five curves fluctuate narrowly around the zero line before the event date, indicating a stable market in which the actual returns of each sector were broadly consistent with the CAPM model's predictions. Around the event date, the cumulative abnormal return of semiconductors (SOXX) fell by approximately 7 percentage points over the $[-1, +1]$ window, with an actual single-day decline of 7.8% on the event date, while communication services (XLC) moved upward against the trend. This contrast indicates that the shock brought by the open-source LLM was highly concentrated in the computing-power layer: the relative price performance of the sectors presents a pattern of "a falling computing-power layer alongside a resilient application layer," which is consistent with a scenario in which capital is reallocated from computing power toward applications within the industry chain. In the following days, SOXX recovered gradually, regaining only about half of its cumulative losses within a week, suggesting that the market progressively repaired itself after the shock. In the right panel (Kimi K3), the decline in the semiconductor sector was shallower and rebounded quickly within five trading days, which may indicate that the market had developed a certain degree of immunity after experiencing the DeepSeek shock.

Table 3. Cumulative abnormal returns around the two open-source releases

Event	Window	SOXX	IGV	XLK	XLC	XLU
DeepSeek-R1 (day 0 = 2025-01-27)	$[-1, +1]$	-7.28	2.00	-1.78	2.62	-2.61
	$[0, +5]$	-2.82	1.28	-1.84	3.56	-1.33
Kimi K3 (day 0 = 2026-07-16)	$[-1, +1]$	-6.28	0.92	-2.69	0.26	-1.05
	$[0, +5]$	3.49	-3.93	1.74	-5.39	2.31




Figure 3. Cumulative abnormal return paths over the ten trading days before and after the two open-source releases

4.4. Long-term changes in risk structure after the DeepSeek shock

Table 4 and Figure 4 report the CAPM parameters estimated separately over the 875 trading days before and the 400 trading days after the DeepSeek-R1 release date. As shown, the advent of open-source LLMs has been accompanied by a structural reshaping of systematic risk across the sectors of the AI industry chain: the computing-power layer became markedly more sensitive to market fluctuations, with the beta of SOXX rising from 1.695 to 2.014, an increase of approximately 19%, while technology (XLK) also rose modestly from 1.346 to 1.452. By contrast, software (IGV) declined from 1.375 to 1.192, and communication services (XLC) fell from 1.100 to 0.797, dropping below the market average, utilities (XLU), as a defensive sector, declined further from 0.503 to 0.319. These contrasting movements indicate that, following the DeepSeek shock, AI-related systematic risk exhibited a tendency to concentrate in semiconductors—the core segment of the industry chain—while the co-movement of the software and platform sectors with the broader market weakened markedly. Re-estimation after excluding the data of the event week leaves the direction and magnitude of these beta changes essentially unchanged, indicating that this shift in risk structure was not driven by the short-term price movements of the event week, but constitutes a structural change that has persisted for more than a year and a half.

Table 4. CAPM parameter estimates before and after the DeepSeek-R1 shock

ETF	β (pre)	α (pre, daily %)	β (post)	α (post, daily %)	$\Delta\beta$
SOXX	1.695	-0.005	2.014	0.117	0.318
IGV	1.375	-0.018	1.192	-0.047	-0.183
XLK	1.346	0.002	1.452	0.035	0.106
XLC	1.100	-0.014	0.797	-0.018	-0.303
XLU	0.503	0.005	0.319	0.001	-0.184




Figure 4. Betas of the five sector ETFs before and after the DeepSeek-R1 shock

5. Conclusion

This study applies the CAPM to examine the impact of two open-source LLM releases (DeepSeek-R1 and Kimi K3) on the returns and risk of U.S. AI-related sector ETFs. Based on the foregoing analysis and discussion, the findings are summarized as follows.

First, both open-source releases imposed notable shocks on the semiconductor sector (SOXX). On the DeepSeek release date, SOXX actually fell by 7.8% and cumulatively underperformed the model prediction by 7.28% over the three trading days surrounding the event, at the release of Kimi K3, SOXX cumulatively underperformed by 6.28% over three days. In both events, the computing-power layer fell substantially behind the model prediction while the application layer remained relatively resilient, presenting a short-term pattern of repricing within the industry chain.

Second, the two shocks differed markedly in their speed of recovery. After the DeepSeek shock, SOXX recovered only about half of its losses within a week, whereas after the Kimi K3 shock, its cumulative abnormal return was positive within five trading days. This difference may indicate that the market had developed a certain degree of immunity following the first shock, responding to similar events more quickly and more mildly.

Third, the risk structure of the sectors underwent persistent changes after the DeepSeek shock. The beta of SOXX rose from 1.70 to 2.01, and this change holds after excluding the event-week data, indicating that systematic risk continued to concentrate in semiconductors throughout the post-event sample period of a year and a half. Meanwhile, the betas of software (IGV), communication services (XLC), and utilities (XLU) all declined markedly, with communication services falling below the market average. This suggests a clear divergence in the market's risk pricing of different segments of the AI industry chain, whose co-movement with the broader market is no longer uniform.

This study has several limitations. First, the sample contains only five U.S. sector ETFs and only two comparable events, so the generalizability of the conclusions awaits verification with larger samples. Second, this study employs only the basic single-factor CAPM without controlling for other factors, such as momentum and size, and the comparison of betas before and after the event is descriptive, without formal statistical significance testing. Third, in the weeks preceding the Kimi K3 event, the semiconductor sector was itself in a highly volatile correction period, and the event window partially overlapped with the U.S. earnings season, therefore, the long-window results for this event should be interpreted with caution.

References

- [1] Ayub, U., Kausar, S., Noreen, U., et al. (2020). Downside risk-based six-factor capital asset pricing model (CAPM): A new paradigm in asset pricing. *Sustainability*, 12: 6756.
- [2] Brusov, P. N., Filatova, T. V., & Kulik, V. L. (2024). Incorporating CAPM into capital structure theories: Accounting for business and financial risks. *Finance: Theory and Practice*, 28: 83–108.
- [3] Habis, H. (2024). A three-period extension of the CAPM. *Journal of Economic Studies*, 51: 200–211.
- [4] Pooja, R., Kayal, P., & Maiti, M. (2024). Enhancing portfolio decision-making: A capital asset pricing model-based clustering analysis. *Journal of Economic Studies*, 51: 358–379.
- [5] Ai, X., & Zhang, Q. (2025). The applicability of the capital asset pricing model: An empirical test. *Frontiers in Economics and Management*, 6: 112–116.
- [6] Hundal, S., Eskola, A., & Tuan, D. (2019). Risk–return relationship in the Finnish stock market in the light of capital asset pricing model (CAPM). *Journal of Transnational Management*, 24: 305–322.
- [7] Wijaya, E., & Ferrari, A. (2020). Stocks investment decision making capital asset pricing model (CAPM). *Jurnal Manajemen*, 24: 93–108.
- [8] Xiao, C. (2022). An empirical test of CAPM before and after the pandemic outbreak. *Proceedings of the 2022 7th International Conference on Financial Innovation and Economic Development (ICFIED 2022)*. Atlantis Press, 2449–2457.
- [9] Cao, H. (2022). The practice and analysis of CAPM on energy commodity ETFs. *Proceedings of the 2022 4th International Conference on Economic Management and Cultural Industry (ICEMCI 2022)*. Atlantis Press, 1074–1084.
- [10] Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance*, 19: 425–442.