

The Release of GPT-4 and Trading Activity in U.S. Equity Markets

Ruohan Qiao*

Shoreline Community College, 16101 Greenwood Ave N, Shoreline, Wa 98133 USA

**Corresponding author: Ruohan Qiao.*

Abstract

This paper analyzes changes in U.S. stock market trading activity around the release of GPT-4 in March 2023. We use a segmented interrupted time-series design with level-discontinuity and regression-kink components. The analysis uses monthly market-wide trading data from March 2020 to May 2026 and supplementary FINRA measures of ATS, non-ATS and total OTC trading. The preferred specification with a 24-month window before and after the event estimates a 22.6% increase in market-wide dollar volume at the event date. The estimate remains positive under multiple robustness checks. A growth-rate ARDL model also finds a significant event-month impulse, but not a sustained increase in subsequent growth. The estimate loses statistical significance once March 2023 is excluded; May and June 2023 are also plausible structural break dates, and the FINRA measures show no clear immediate discontinuity. The main result is therefore a positive but short-lived increase in trading activity that coincides with the release of GPT-4. The estimates do not establish that GPT-4 caused the increase. The modeled increase in volume reflects repeated turnover of identical securities, not new wealth, investment returns or real output.

Keywords

GPT-4, investor attention, interrupted time series, regression discontinuity, regression kink

1. Introduction

The release of generative artificial intelligence (AI) has raised questions about its effects on financial markets through individual productivity, investor behavior and trading activity. On 14 March 2023, OpenAI launched GPT-4, an improvement over its predecessor on complex reasoning, understanding text and performing specialized tasks [1]. These tools allow investors to understand corporate disclosures, interpret news, write code and perform scenario analysis at a lower cost. While prior evidence documents productivity gains from generative AI across professional writing, customer service and software engineering [2-4], such studies do not examine whether generative AI alters aggregate market trading activity.

Recently, large language models are used as information-processing technologies. Lopez-Lira and Tang [5] demonstrate that GPT-based evaluations of news headlines are predictive of future stock returns. The effect

on trading volume is theoretically ambiguous. Lower participation and research costs may increase entry and turnover, while similar model outputs may reduce disagreement and, in some settings, trading.

This paper treats the public release of GPT-4 as a dated market event and asks three questions. First, was there an immediate change in level for overall market dollar volume, share volume, or trade counts? Second, was the post-release growth path different from the pre-event trajectory? Third, was the change in activity a redistribution between ATS and non-ATS venues or a more general shift across OTC channels?

In order to answer these questions, the paper makes four contributions. Regarding the research design, it separates an immediate level shift from a post-event trend change, thus avoiding the identification problems associated with using a single post-event dummy. In terms of data, it combines Cboe market-wide aggregates with FINRA OTC transparency data in order to compare trading changes across the different market segments [6,7]. With respect to identification, it uses event-time coefficients, growth-rate ARDL models, alternative breakpoints, rolling windows, AR/ARMA error structures, robust regression, leave-one-month-out procedures, and negative control tests to check the results for sensitivity to specification, timing, and outliers. Finally, in economic interpretation, it converts the log-model coefficients into amounts in dollar terms and stresses that any estimated increase represents repeated turnover of securities and not the creation of new wealth or real output.

2. Conceptual Framework and Hypotheses

2.1 Information-Processing Costs

GPT-4 can summarize lengthy documents, explain technical concepts and write computer code. These capabilities reduce the fixed costs of processing financial information and allow market participants with limited time or limited technical background to make initial assessments. This implies an increase in both trade count and dollar volume. If adoption is slow the market response may look like a change in the slope of the growth path. If the release precipitates a concentrated burst of attention the response may look like a one-off jump in levels.

Hypothesis 1. The release of GPT-4 is associated with an immediate increase in trading activity in the U.S. stock market.

Hypothesis 2. After the release of GPT-4, the growth path of trading activity is steeper than what the pre-release trend would predict.

2.2 Attention, Divergence, and the Choice of Trading Venues

Limited attention theories suggest that investors cannot process all available information simultaneously. Da, Engelberg, and Gao [8] show that Internet search activity proxies retail attention and correlates with short-run price pressure. Barber and Odean [9] find that individual investors tend to buy attention-capturing stocks. The GPT-4 release could affect trading through three channels: focus on the AI theme, repricing of technology supply chains, and expectations that AI tools will ease investment research. Greater attention and diverging views could increase turnover and options activity. The analysis uses Google Trends data and AI-supply-chain assets as tentative indicators.

FINRA data offer a separate look at the composition of trading venues. The ATS market share could grow if the new technology is taken up mostly by institutional investors or automated strategies. If instead order generation expands to include broker internalization and other OTC channels, both ATS and non-ATS volumes could grow in parallel with the ATS share remaining roughly stable.

Hypothesis 3. If the main effect of GPT-4 is on venue allocation, the ATS share of total OTC notional volume should change detectably after the release.

3. Data

3.1 Sources and Sample Construction

The full market sample is 75 months long, from March 2020 to May 2026. Total U.S. equity share volume, dollar volume and trade counts are from Cboe Historical Market Volume. The S&P 500 and Nasdaq Composite

series are from FRED; QQQ and AI-supply-chain asset histories are from the official historical interface of Nasdaq. VIX is from FRED, and the National Financial Conditions Index (NFCI) and Adjusted NFCI are from the Federal Reserve Bank of Chicago. The federal funds rate, 10-year Treasury yield, unemployment rate, Consumer Price Index and industrial production are from FRED.

The FINRA supplement covers March 2022 through March 2024. Official weekly OTC Transparency records are aggregated according to the month containing week Start Date. The aggregation retains ATS and non-ATS observations separately and defines total OTC activity as their sum. Rows with missing product Type Code are excluded because they are cross-product totals that would duplicate CTS, UTP, and OTCE components. Each weekly date retains six venue-by-reporting-system cells. The final monthly data contain no duplicated months or missing notional amounts, and the ATS-plus-non-ATS identity holds exactly.

Table 1: Data sources and construction

Module	Variables	Source	Frequency and processing
Market-wide trading	Dollar volume, shares, trades	Cboe Historical Market Volume	Aggregated to monthly frequency; Mar. 2020–May. 2026
OTC trading	ATS, non-ATS, total OTC	FINRA OTC Transparency weeklySummary API	Weekly CTS/UTP/OTCE records aggregated by month; Mar. 2022–Mar. 2024
Market prices	S&P 500, Nasdaq, QQQ, AI assets	FRED and Nasdaq official history	Month-end prices, returns, and volume
Financial conditions	VIX, NFCI, ANFCI	FRED and Federal Reserve Bank of Chicago	Daily or weekly data aggregated to monthly frequency
Macroeconomic controls	Rates, unemployment, CPI, industrial production	FRED	Monthly series or daily data aggregated to monthly frequency
AI attention	ChatGPT, generative AI, LLM, AI chatbot	Google Trends	Exploratory only; no common-anchor scaling

3.2 Variable Construction

The primary trading outcomes enter in natural logarithms. The monthly event indicator equals one in March 2023 and thereafter. The variable *event_time* counts the number of months from March 2023. The variable *post_slope* is zero before the event and increases by one in each month thereafter. The event indicator measures the immediate level shift; *post_slope* measures the incremental monthly slope.

The preferred controls are the monthly mean of VIX, the monthly S&P 500 return, the monthly change in the federal funds rate, and the monthly mean of NFCI. Seasonality is modeled using annual sine and cosine terms.

FINRA specifications also control for the log number of weekly reporting periods assigned to each month. Seasonal terms are not used to save degrees of freedom due to the small size of the FINRA sample (25 observations).


3.3 Descriptive Evidence

Table 2: Descriptive statistics

Variable	N	Mean	Std. dev.	Median	Minimum	Maximum
Federal funds rate	75	2.844	2.187	3.72	0.05	5.33
FINRA total OTC notional (USD)	25	4,627,699,734,889	730,083,337,348	4,461,424,257,824	3,541,762,200,665	6,004,262,228,509
NFCI	75	-0.4109	0.1782	-0.4378	-0.6907	0.2462
S&P 500 monthly return	74	0.01571	0.04666	0.02194	-0.0934	0.1268
Market-wide share volume	75	273,094,917,371	66,098,971,426	242,621,744,025	193,808,198,661	482,907,491,977
Market-wide trade count	75	1,709,844,108	465,129,107	1,564,225,612	1,075,030,056	3,134,973,893
Market-wide dollar volume (USD)	75	13,149,208,817,789	3,510,758,872,953	11,956,227,180,041	8,511,996,771,785	24,240,850,561,849
VIX monthly mean	75	20.98	7.017	19.76	12.67	57.74

The three market-wide trading measures in Figure 1 are set to 100 in the first month of the sample. Trading was high in the early days of the pandemic, then dipped, then bounced back. The large cycle shows that a simple before-after comparison would not isolate an effect of GPT-4.

Figure 1: Market-wide U.S. equity trading activity (March 2020 = 100)



4. Empirical Strategy

4.1 Segmented Interrupted Time Series

The baseline specification is given by

$$\ln(Y_t) = \alpha + \beta_1 \text{Time}_t + \beta_2 \text{Post}_t + \beta_3 \text{PostTrend}_t + \gamma' X_t + s_t + \varepsilon_t, \quad (1)$$

where Y_t denotes dollar trading volume, share volume, or the number of trades. The trend variable Time_t is linear and equals zero in March 2023. The indicator Post_t equals one from March 2023 onward. The term $\text{PostTrend}_t = \max(\text{Time}_t, 0)$ is the post-event trend. The vector X_t includes market and macroeconomic controls, and s_t accounts for seasonality.

In this specification, β_2 measures the conditional change in the level of the outcome at the event date, and β_3 measures the change in the subsequent trend relative to the pre-event period. Statistical inference is based on Newey–West heteroskedasticity- and autocorrelation-consistent standard errors, using a maximum lag length of three months [10].

All hypothesis tests are two-sided. For models with logarithmic dependent variables, estimated coefficients are converted into percentage effects using $100[\exp(\hat{\beta}) - 1]$, rather than the linear approximation $100\hat{\beta}$.

A specification with only Post_t captures an immediate level discontinuity around the date of the event. The addition of PostTrend_t also provides for a different slope after the event than the pre-event trend, which can be interpreted as a regression-kink component. The design is better characterized as a segmented interrupted time series rather than a traditional regression discontinuity design in that there is no assignment rule that generates treatment around a cutoff. The time threshold alone is not enough to identify causality, since there may be contemporaneous market shocks in the month of release.

4.2 Model Selection and Economic Magnitudes

We compare four specifications for each outcome: no break, jump only, slope only, and jump plus slope. The comparison is done based on AIC, BIC and adjusted R^2 . The formal tests for the discontinuity and the change in slope are given by the full model. The dollar counterfactual magnitudes are calculated using the more parsimonious form preferred by the BIC, which sets the post-event terms to zero but retains the realized controls and seasonality. Predictions are back-transformed from logs with a residual-based correction. Ninety-five percent parameter intervals are calculated from 2,000 samples of the coefficient distribution using the Newey–West covariance matrix.

4.3 Identification Diagnostics and Robustness Design

We assume the pre-event path approximates the counterfactual absent the event [11]. This assumption is dealt with by five sets of checks.

First, we vary the window (± 12 , ± 18 , and ± 24 months), market-return control, policy-rate treatment, and trend function, and compare local with full-sample estimates. Second, identical ± 18 -month models are re-

estimated for candidate breakpoints from December 2022 through June 2023. March 2023 is excluded and implement a donut design that omits February through April 2023.

Third, a grouped event study uses the month immediately before the event as the reference and groups the surrounding 24 months into three-month bins. A Newey–West joint F-test evaluates the four pre-event groups.

$$\Delta \ln(Y_t) = \alpha + \theta_0 \text{Pulse}_t + \theta_1 \text{PostFull}_t + \phi \Delta \ln(Y_{t-1}) + \gamma' \Delta X_t + s_t + u_t, \quad (2)$$

where Pulse_t equals one only in March 2023 and PostFull_t equals one beginning in April 2023. This ARDL(1) model specification distinguishes the impulse effect of the “event month” from persistent differences in the average monthly growth rate and reduces the model’s sensitivity to non-stationarity in the levels of the outcome variables.

Fourth, robustness checks are conducted by varying the Newey–West lag order between 0 and 12, estimating GLS models with AR(1), AR(2), and ARMA(1,1) error terms, employing Huber M-estimation, and performing a month-by-month sample exclusion analysis.

Finally, the S&P 500 return, changes in the VIX index, changes in the unemployment rate, and the industrial production growth rate were used as negative control outcome variables. Supplementary tests were also conducted using the ChatGPT release date, placebo dates, and an exploratory GPT-4/GPT-5 two-break model specification.

5. Results

5.1 Market-Wide Level and Slope Changes


Table 3 presents the preferred estimation results using a ± 24 -month window. At the threshold, transaction value is 22.6% higher ($p = 0.038$), the number of transactions is 22.8% higher ($p = 0.001$), and trading volume (number of shares) is 16.7% higher ($p = 0.008$). The slope coefficients for transaction value and the number of transactions following the event are not statistically significant. The estimated change in the slope for trading volume is approximately 2.4% per month, though this result is not robust under the dynamic error specification.

Table 3: Segmented interrupted time-series estimates, ± 24 -month window

Outcome	Effect	Estimate	HAC SE	p-value	BH-adjusted p	N
Share volume	Immediate level difference	16.65%	0.055	0.008	0.012	49
Share volume	Monthly slope change	2.41%	0.009	0.014	0.043	49
Dollar volume	Immediate level difference	22.56%	0.095	0.038	0.038	49
Dollar volume	Monthly slope change	0.81%	0.009	0.384	0.384	49
Trade count	Immediate level difference	22.81%	0.059	0.001	0.004	49
Trade count	Monthly slope change	1.14%	0.007	0.135	0.203	49

Figure 2 shows the observed U.S. dollar trading volume, the fitted event model, and the no-break counterfactual. The fitted path is above the counterfactual after March 2023 most of the time, but the confidence intervals are wide, indicating large parameter uncertainty.

Figure 2: Market-wide dollar volume: fitted path and no-break counterfactual



A separate extension plots covariate-adjusted monthly observations and fitted trends on each side of March 2023, but does not replace the baseline specification.

Figure 3: Covariate-adjusted local segmented regression around March 2023




Figure 3 makes the fitted break visible. The covariate-adjusted segmented model estimates an immediate level difference of 22.6% ($p=0.038$), while the estimated monthly slope change is 0.8% ($p=0.384$). The separate growth-rate model estimates a 26.5% event-month impulse ($p<0.001$), but the average growth difference after March is only 0.8% per month ($p=0.771$). The results are positive at the event date, but they fit a short-run impulse better than a persistent acceleration in trading growth.

5.2 Model-Implied Economic Magnitudes

Over the first 12 months after March 2023, the model implies cumulative nominal trading volume of approximately USD 24.5 trillion above the no-break counterfactual. The corresponding 95% parameter interval ranges from USD 1.6 trillion to USD 45.5 trillion. This gap is equivalent to about 22.3% of the counterfactual trading volume. Valued at end-of-sample prices using the CPI, the estimated real-dollar gap is approximately USD 27.4 trillion.

Table 4: Model-implied cumulative dollar-volume gaps

Market	BIC-preferred model	12-month gap (USD tn)	95% parameter interval	Gap / counterfactual
Market-wide (Cboe)	Jump only	24.5	[1.6, 45.5]	22.3%
FINRA total OTC	Slope only	10.5	[5.9, 14.7]	22.7%

5.3 Finra Otc Evidence

The FINRA results tell a different story. The difference in total OTC notional volume at the event date is 0.8% ($p = 0.823$); the GLS AR(1) estimate is 3.1% ($p = 0.696$); and the specification excluding March 2023 is -0.2% ($p = 0.971$). The FINRA sample does not show evidence of an immediate OTC discontinuity.


The post-event slope is 3.75% per month ($p < 0.001$) in the ± 12 -month model, 4.00% under AR(1) errors ($p = 0.002$), and 4.18% when April 2023 is the first fully treated month ($p < 0.001$). BIC favors slope-only models for ATS, non-ATS and total OTC notional volume. The slope estimate drops to 1.79% ($p = 0.341$) in the ± 6 -month window and 1.33% ($p = 0.132$) in the ± 9 -month window. The p-value for the placebo in the small sample is around 0.667. The widest FINRA window is consistent with a shift in the growth path, but the evidence is not strong enough to make a causal interpretation.

Table 5: FINRA total OTC robustness results

Effect	OLS estimate	OLS p	GLS AR(1)	GLS p	Full-month estimate	Full-month p	Assessment
Immediate level difference	0.83%	0.823	3.09%	0.696	-0.18%	0.971	No stable level discontinuity
Monthly slope change	3.75%	<0.001	4.00%	0.002	4.18%	<0.001	Supported only in the widest window

Figure 4 shows the slope-only FINRA specification. The 12-month cumulative gap from this model is about USD 10.5 trillion (95% parameter interval: USD 5.9–14.7 trillion) or 22.7% of the model counterfactual. The gap arises from a difference in slopes that builds up over time, making the estimate sensitive to the choice of window and to how far the trend is extrapolated.

Figure 4: FINRA total OTC notional volume: fitted path and no-break counterfactual



The estimated change in notional trading volume was 6.7% ($p = 0.184$) for ATS and -1.7% ($p = 0.618$) for non-ATS. The average ATS share of total OTC notional volume was 0.71 percentage points higher after the event, with a 95% confidence interval of -1.07 to 2.47. The confidence interval includes zero, indicating that the data show no meaningful shift in execution activity from non-ATS to ATS venues.

6. Robustness and Alternative Explanations

6.1 Error Structure and Influential Observations

The estimated break in market-wide dollar volume remains positive for the ± 12 -, ± 18 - and ± 24 -month windows with point estimates in the range from 22.6% to 26.9%. Of the three, two make the 5% significance bar. Varying the Newey–West lag between 0 and 12 produces p -values between 0.006 and 0.044. The estimated discontinuities under AR(1), AR(2) and ARMA(1,1) error models are 26.5%, 30.1% and 28.6%, and all are below the 5% level. The estimate from Huber M-estimation is 29.6% ($p < 0.001$).

All 49 leave-one-month-out estimates are positive with a range of 17.2% to 30.3%, and 43 of 49 are significant at the 5% level. These tests show that the local discontinuity is not an artifact of a particular HAC lag, a particular error model, or one extreme month. The full-sample estimate is -7.0% ($p = 0.519$), and the local pattern does not hold for the whole 2020–2026 period.

The post-event slope estimate depends on the trend specification. Under the linear segmented model the change of the slope is 0.81% ($p = 0.384$). Adding a common quadratic trend increases this to 5.72% ($p < 0.001$). The quadratic specification yields lower AIC and BIC values, but also increases the condition number of the design matrix, which means more collinearity. The presence of a persistent post-event slope change is a function of the assumed functional form. The differenced ARDL and the linear segmented model give the more conservative reading.

Table 6: Identification and robustness audit for market-wide dollar volume


Diagnostic	Result	Implication
HAC lag	$p = 0.006\text{--}0.044$ for lags 0–12	Inference does not depend on one HAC lag
AR/ARMA errors	26.5%–30.1%; all $p < 0.05$	Positive under alternative serial-correlation models
Huber M-estimation	29.6%; $p < 0.001$	Not driven by conventional OLS residual weighting
Leave one month out	17.2%–30.3%; 43/49 $p < 0.05$	Direction is stable; significance is partly month-sensitive
Growth-rate ARDL	Impulse 26.5%; later growth 0.78%	One-time impulse, not persistent growth acceleration
Pre-event joint test	F-test $p = 0.409$	No detected pre-event differences; low power remains
Candidate breakpoints	March, May, and June 2023 significant	Exact GPT-4 timing is not unique
Exclude March 2023	17.0%; $p = 0.127$	Partial-month treatment affects inference
Donut window	30.3%; $p = 0.047$	Spring 2023 change remains, but timing shifts
Negative controls	All immediate terms BH $p > 0.05$	No broad simultaneous macro discontinuity
Full sample	-7.0%; $p = 0.519$	Local result does not extrapolate

6.2 Dynamic Path and Growth-Rate Evidence

The grouped event study does not reject the joint null for the four pre-event periods ($p = 0.409$), though this test has low power. Dollar volume in the event month is 47.9% higher than the month immediately preceding the event ($p = 0.005$), while none of the four post-event groups are individually significant. Figure 6.1 fits better with a concentrated impulse than with a gradually broadening effect.

The growth-rate ARDL makes the point sharper. The event-month impulse in dollar-volume growth is 26.5% ($p < 0.001$), while the average difference from April 2023 onward is 0.78% per month ($p = 0.771$). Share volume and trade count show similar impulses over one month but no lasting growth difference. These models favor a short-run shift close to the event over a permanent change in the growth rate of trading.


Figure 5: Grouped event-study estimates for market-wide dollar volume



6.3 Breakpoint Timing, Placebos, and Negative Controls

The main problem with causal attribution is timing sensitivity. For the same ± 18 -month specifications, the March 2023 estimate is 24.3% ($p = 0.021$) while the May and June 2023 estimates are 27.9% ($p < 0.001$) and 23.2% ($p = 0.004$). If we drop March and use April as the first full post-event month we get 17.0% ($p = 0.127$). Dropping February through April and treating May as the first post-event month gives 30.3% ($p = 0.047$). Figure 6.2 supports a local change during spring 2023 but does not pin it to the month of the GPT-4 release.

Figure 6: Sensitivity of the dollar-volume discontinuity to the candidate breakpoint



The empirical two-sided p -value of the placebo-date test is about 0.077. None of the immediate effects for the four negative-control outcomes cross the Benjamini–Hochberg threshold. This reduces the worry about a broad macroeconomic break on the same date, though the unadjusted post-event slope of S&P 500 returns is significant, which suggests the technology-stock rally may have altered the post-event environment. The dynamic and error-structure checks are consistent with a local change, while the candidate-date, full-sample, and event-month-exclusion results undermine the case for GPT-4 as the cause.

7. Discussion and Economic Implications

The increase in trading activity around March 2023 cannot be taken as a direct improvement in market efficiency. Higher volume could be a sign of faster incorporation of public news into prices, or of fleeting attention, speculative turnover or repeated trading of the same securities. The empirical estimates do not separate these channels.

But the fact that the volume spike was transient also matters. The estimated change is concentrated in the month of the event rather than coming with a sustained shift. When we exclude observations from March 2023, the estimated discontinuity loses statistical significance and May and June are plausible alternative break dates. The change in trading activity appears to be more of a transitory episode than a permanent change in the price discovery process.

Stock and option dollar volumes moved together over the same period, consistent with investors using options alongside equities to express views or hedge amid elevated market interest. This interpretation is tentative: the option proxy is from Cboe monthly series, not comprehensive OCC data, and both markets may respond to common volatility or technology-sector conditions.

FINRA data show no detectable shift in ATS market share; ATS and non-ATS volumes moved in the same direction, implying an economy-wide expansion in order flow rather than a compositional shift in execution venues.

The estimated dollar gap is a measure of trading volume; it does not represent new wealth, new production, capital formation or profits for investors. Cumulative volume can be large even when there is little change in the underlying value of listed companies, as the same securities are traded many times. The estimated change bears on market operations: increasing volume affects demand for liquidity, transaction costs, market-making capacity and the load on trading infrastructure. The impact on market quality depends on whether the additional volume is the result of informed trading or uninformed turnover.

8. Conclusion

Using a segmented interrupted time-series model, we find an approximate 22–23% increase in dollar volume and trade count at the event date under the local specification. This increase is positive across a range of robustness checks—varying HAC lag length, substituting AR and ARMA error structures, applying Huber M-estimation, and dropping months individually—and the ARDL model shows a sharp spike in volume growth during the event month, though no sustained increase in the average growth rate follows. These estimates do not imply causation: the full-sample model does not capture the local discontinuity; the coefficient loses significance when March 2023 is excluded; May and June 2023 are also candidate breakpoints; FINRA data show no immediate OTC shift, with the positive post-event slope significant only in the widest window; and ATS market share remains unchanged.

The evidence across all specifications points to a localized shift in U.S. stock trading activity in the spring of 2023. This change coincides with the early diffusion of generative AI, but the data do not allow a causal link to GPT-4. To show such a link, higher-frequency trading data, a valid control group and direct measures of generative AI adoption would be needed. The one-breakpoint approach does not exclude other shocks around March 2023. The May and June 2023 data also indicate potential breakpoints and the monthly data alone cannot determine March as the only date. The study lacks a control group unexposed to generative AI. Future work could improve causal identification using controlled interrupted time-series designs, the synthetic control method, or cross-country comparisons that exploit variation in timing and intensity of technology adoption.

The monthly frequency of the data limits the precision of the event-time analysis. March 2023 data contains pre-release and post-release dates, so the month is only partially treated. Data at the daily or weekly level would help separate the immediate response to the launch of GPT-4 from the broader technology rally in spring 2023 and the more gradual adoption of generative AI tools.

The FINRA sample is restricted to the 12 months before and after the event. The estimated post-event slope is significant only in the widest available window, so the cumulative dollar trading-volume gap is sensitive to the sample period and the extrapolation horizon. The Google Trends series are not normalized to a common baseline and the options variables do not cover the entire OCC market.

The aggregate market data do not distinguish between types of investors, directions of trades, or differences among firms. Future studies could use CAT, TAQ, 13F filings, retail order designations, or firm-level trading records to examine whether market responses are more pronounced for retail investors, AI-related firms, small-cap stocks, or firms with high information asymmetry.

References

- [1] OpenAI. (2023). GPT-4 technical report. <https://cdn.openai.com/papers/gpt-4.pdf>
- [2] Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192. <https://doi.org/10.1126/science.adh2586>
- [3] Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *Quarterly Journal of Economics*, 140(2), 889–942. <https://doi.org/10.1093/qje/qjae044>
- [4] Peng, S., Kalliamvakou, E., Cihon, P., & Demirer, M. (2023). The impact of AI on developer productivity: Evidence from GitHub Copilot. arXiv:2302.06590. <https://arxiv.org/abs/2302.06590>
- [5] Lopez-Lira, A., & Tang, Y. (2023). Can ChatGPT forecast stock price movements? Return predictability and large language models. SSRN Working Paper 4412788. <https://doi.org/10.2139/ssrn.4412788>
- [6] Cboe Global Markets. Historical market volume. <https://www.cboe.com/markets/us/equities/market-statistics/historical-market-volume>
- [7] Financial Industry Regulatory Authority. OTC transparency data. <https://www.finra.org/filing-reporting/otc-transparency>
- [8] Da, Z., Engelberg, J., & Gao, P. (2011). In search of attention. *Journal of Finance*, 66(5), 1461–1499. <https://doi.org/10.1111/j.1540-6261.2011.01679.x>
- [9] Barber, B. M., & Odean, T. (2008). All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. *Review of Financial Studies*, 21(2), 785–818. <https://doi.org/10.1093/rfs/hhm079>
- [10] Newey, W. K., & West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3), 703–708.
- [11] Bernal, J. L., Cummins, S., & Gasparrini, A. (2017). Interrupted time series regression for the evaluation of public health interventions: A tutorial. *International Journal of Epidemiology*, 46(1), 348–355. <https://doi.org/10.1093/ije/dyw098>

Funding

This research received no external funding.

Conflicts of Interest

The authors declare no conflict of interest.

Acknowledgment

This paper is an output of the science project.

Open Access

This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

