

The Relationship Between NVIDIA Stock Price Changes and Nasdaq-100 Market Volatility During the AI Boom: Evidence from 2023–2025

Yihan Sun*

Beijing Huijia Private School,
Beijing, 102299, China

*Corresponding author:
27sunyihan@huijia.edu.cn

Abstract:

The rapid growth of generative artificial intelligence made NVIDIA one of the most closely watched companies in the U.S. stock market. This paper studies the relationship between NVIDIA's daily price changes and Nasdaq-100 market volatility during 2023–2025. Daily log returns were calculated from 752 matched closing-price observations, leaving 751 return observations. Absolute daily return is treated as an observable proxy for short-term volatility rather than as the market's unobservable conditional variance. The analysis combines descriptive statistics, correlation, a contemporaneous regression, and a fixed-parameter one-day-ahead forecast test. Inference uses Newey–West heteroskedasticity-and-autocorrelation-consistent (HAC) standard errors with five lags. NVIDIA and Nasdaq-100 absolute returns had a correlation of 0.583. In the same-day regression, a one-percentage-point increase in NVIDIA absolute return was associated with a 0.237-percentage-point increase in Nasdaq-100 absolute return (HAC SE = 0.044; 95% CI = [0.150, 0.323]). The forecasting evidence was weak: adding previous-day NVIDIA volatility reduced 2025 MAE by 0.4% but increased RMSE by 0.4%. HAC loss-differential tests did not reject equal predictive accuracy under absolute loss ($p=0.342$) or squared loss ($p=0.300$). The results therefore support a strong same-day association, not a causal effect or a reliable one-day-ahead forecasting advantage.

Keywords: NVIDIA; Nasdaq-100; artificial intelligence; daily returns; volatility.

1. Introduction

1.1 Background and Research Phenomenon

Generative artificial intelligence expanded rapidly after the public release of tools such as ChatGPT. Businesses increased investment in AI systems and the graphics processing units that support them. NVIDIA designs GPUs and related computing platforms, making it one of the most visible companies associated with the AI boom. Its fiscal 2024 annual report identified generative AI as an important driver of accelerated-computing demand, while its fiscal 2025 report documented continued expansion by companies and cloud providers [1, 2]. Together, the reports support treating January 2023 through December 2025 as one coherent AI-boom research window rather than as unrelated calendar years.

The same period was important for the stock market. NVIDIA's market value increased greatly, and the company became a major member of the Nasdaq-100 Index. The index measures 100 of the largest Nasdaq-listed non-financial companies and uses modified market-capitalization weights. Its company-level concentration rules operate in two stages: if any initial company weight exceeds 24%, weights are adjusted so that no company exceeds 20%; if companies above 4.5% would together reach at least 48%, that group is reduced to 40% [3]. Thus, a large company can influence the index, but it cannot determine the entire index by itself. NVIDIA's dual role as a large constituent and a symbol of AI investment creates a plausible same-day connection with broader technology-share movements.

1.2 Research Objective and Questions

Focusing on 2023–2025, this paper investigates two questions: the contemporaneous linkage between NVIDIA returns and Nasdaq-100 volatility, and the predictive power of lagged NVIDIA volatility for next-day market moves. Using 751 daily observations (split at 2025 for out-of-sample testing), the study identifies statistical associations without asserting causality, maintaining a narrow scope to ensure robust, defensible conclusions.

1.3 Academic and Practical Significance

Academically, the paper provides a focused case study of firm-index volatility linkage during an unusual technology cycle. It separates same-day comovement from lagged predictive content and shows how a transparent daily proxy can be used when intraday data are unavailable. Practically, the analysis discourages treating a large NVIDIA movement as an automatic signal for the next trading day. Same-day comovement may be strong even when forecasting gains are too small to guide a decision.

1.4 Paper Structure

The remainder of the paper reviews volatility measurement, forecasting, information transmission, and the research gap. It then states the hypotheses, defines the data and variables, explains the empirical design, reports baseline and robustness results, discusses limitations, and concludes.

2. Literature Review

2.1 Meaning and Measurement of Volatility

Volatility describes the size of price changes, not their direction. A stock that rises by 5 percent and a stock that falls by 5 percent experience different returns, but both experience a large price movement. Researchers have long studied why stock-market volatility changes over time. Schwert examined historical U.S. data and showed that market volatility varies substantially rather than remaining constant [4]. This matters because risk is not equally high on every day or in every year.

Different measures serve different purposes. Intraday realized volatility can combine many within-day returns, but it requires high-frequency data. Rolling standard deviations require a chosen window and lose observations at the beginning of the sample. A GARCH conditional variance depends on a fitted model. Recent research uses realized volatility and large panels of predictors to improve forecasts [5, 6]. This paper instead uses absolute daily return because it is observable, remains in percentage-point units, and can be reproduced directly from public closing prices. Research by Ding, Granger, and Engle also showed that absolute returns contain information about persistent market movements [7]. The proxy is noisy and is not the same as an unobservable conditional variance, so conclusions are stated in terms of an absolute-return volatility proxy.

2.2 Time-Varying Volatility and Classical Models

More advanced financial studies often model changing volatility with ARCH and GARCH methods. Engle introduced the autoregressive conditional heteroscedasticity, or ARCH, model, which allows today's variance to depend on earlier forecast errors [8]. Bollerslev later developed the generalized ARCH, or GARCH, model, which also allows variance to depend on its own past values [9]. These models are important in professional finance, but they require more mathematical and statistical assumptions than are necessary for the present research question.

The model choice is based on the parameter of interest, not merely on simplicity. A univariate ARCH or GARCH model describes one series' own conditional variance;

it does not directly estimate the same-day or lagged cross-series coefficient that this paper asks about. Multivariate GARCH models can study changing covariance, but they require a much larger set of dynamic assumptions. A transparent bivariate regression therefore matches the two research questions more directly. Recent studies also evaluate volatility-timing strategies and the economic value of implied-volatility forecasts, reinforcing the need for out-of-sample evaluation [5, 10, 11, 12].

2.3 Cross-Asset Linkage and Forecasting

Forecasting volatility is more difficult than describing it. Poon and Granger reviewed a large range of volatility-forecasting methods and found that model performance depends on the data, forecast period, and evaluation measure [13]. A model can fit historical data but perform poorly when it is applied to a later period. For this reason, the present paper separates the sample into a training period and a test period. The model is estimated with 2023–2024 data, while 2025 is reserved for evaluating forecasts. This is more convincing than reporting only how well a model explains the same data used to build it.

Research on connectedness also shows that volatility can be shared across markets because assets respond to common information and investor risk preferences [14]. However, connectedness within an index has a special mechanical component: an index constituent contributes directly to the index return. This paper therefore evaluates both same-day linkage and next-day forecasting rather than treating either one as proof of the other.

2.4 Market Efficiency and Information Transmission

The forecasting question is also connected to market efficiency. Fama's review of efficient capital markets explains the idea that available information tends to be incorporated into prices [15]. If investors react quickly to NVIDIA news, much of the related market movement may occur on the same trading day. Yesterday's NVIDIA movement may then have limited value for predicting tomorrow's Nasdaq-100 movement. This does not mean markets are perfectly predictable or perfectly efficient. It simply gives a reason to expect a stronger same-day relationship than a one-day-ahead forecasting relationship.

2.5 Endogeneity and Mechanical Index Linkage

There are also two reasons to be careful when interpreting a relationship between NVIDIA and the Nasdaq-100. First, both can respond to common information. For example, changing expectations about interest rates or technology demand can affect NVIDIA and many other index members together. Second, NVIDIA is itself part of the Nasdaq-100. A portion of the observed relationship is

therefore mechanical: when a large component changes in price, it directly contributes to the index movement. The Nasdaq-100's modified market-capitalization methodology makes this connection important, while its weight-limiting rules prevent the index from becoming a measure of only one company [3].

Recent firm-level research examines AI innovation and stock-price crash risk, while market research studies how investor sentiment fluctuations relate to liquidity and volatility [16, 17]. These possible information channels strengthen the case for interpreting the regressions as associations rather than causal effects.

2.6 Literature Assessment and Research Gap

The reviewed literature leaves three practical gaps for this case study. First, classical volatility research often models an asset's own dynamics, while recent connectedness studies usually compare broad markets rather than one leading company with its own index. Second, recent forecasting studies commonly use intraday data or high-dimensional predictors, leaving room for a transparent design based on public daily prices. Third, same-day association and next-day predictive value are economically different but can be confused in informal discussion of a prominent stock.

This paper addresses those gaps by examining NVIDIA and the Nasdaq-100 during the 2023–2025 AI boom, estimating contemporaneous and one-day-lag relationships in one framework, and evaluating the lagged model on a separate year. The expected same-day relationship is positive, but a strong correlation is not treated as causation or as evidence of reliable forecasting.

3. Research Questions and Hypotheses

The empirical analysis separates contemporaneous association from predictive content.

RQ1: Is NVIDIA's absolute daily return positively associated with the Nasdaq-100's absolute daily return on the same trading day?

H1: In the same-day model, β is evaluated using the coefficient sign, a 95% HAC confidence interval, and the stability of the estimate across years and after 1% upper-tail winsorization. Because both variables are measured on day t , support for H1 indicates association only; it does not identify a causal effect.

RQ2: Does previous-day NVIDIA absolute return improve a one-day-ahead forecast of Nasdaq-100 absolute return beyond information in the index's own previous-day absolute return?

H2: The lagged NVIDIA coefficient satisfies $\beta > 0$ and the extended model produces lower out-of-sample loss than the baseline model. H2 is treated as supported only if the 2025 error reduction is economically meaningful and the mean loss differential is statistically distinguishable from

zero. A significant training-sample coefficient by itself is not sufficient evidence of predictive improvement.

4. Data, Variables, and Descriptive Evidence

4.1 Data Source and Sample

The sample covers January 3, 2023 through December 31, 2025. Historical daily data for NVIDIA (ticker NVDA) and the Nasdaq-100 Index (ticker NDX) were retrieved through the public interface, which reports Nasdaq as the upstream provider [18]. The download contained date, open, high, low, close, and volume fields for NVIDIA, and

the corresponding index-price fields for the Nasdaq-100. Only the date and closing price were needed for this study. Both series contained 752 price observations. Their trading dates matched exactly, so no interpolation was needed. After calculating returns, the first price observation was used only as the starting value, leaving 751 daily return observations from January 4, 2023 through December 31, 2025. The NVIDIA price series was checked around the company's ten-for-one stock split in June 2024. The prices were continuous on a post-split basis, so the split did not create a false daily return.

Table 1 summarizes the sample construction and shows how the 752 price observations become 751 return observations and 750 lagged observations for the forecasting design.

Table 1. Sample construction and estimation windows

Stage	Observations	Use in the analysis
Raw NVIDIA daily closes	752	January 3, 2023–December 31, 2025
Raw Nasdaq-100 daily closes	752	Same calendar as NVIDIA; no interpolation
Matched closing-price pairs	752	Input for return calculation
Matched daily returns	751	Same-day descriptive and regression sample
Observations with one-day lags	750	Available for forecasting design
Training sample	500	2023–2024 model estimation
Test sample	250	2025 out-of-sample evaluation

4.2 Variables

Daily logarithmic return was calculated as:

$$R_t = 100 \ln \left(\frac{P_t}{P_{t-1}} \right) \quad (1)$$

where R_t denotes the daily logarithmic return (in percentage) on trading day t , P_t is the adjusted closing price on day t , P_{t-1} is the adjusted closing price on the previous trading day, and $\ln(\cdot)$ represents the natural logarithm. Multiplying by 100 converts the logarithmic return into percentage points, making the results easier to interpret. The multiplication by 100 expresses the result as a percentage. Logarithmic returns are commonly used because returns from different days can be added more easily than

simple percentage returns. For the small daily changes in this dataset, log returns and ordinary percentage returns are usually close.

Daily volatility was measured as the absolute value of the daily return:

$$V_t = |R_t| \quad (2)$$

where V_t denotes the daily volatility proxy on day t , measured as the absolute value of the daily logarithmic return. The absolute value operator $|\cdot|$ removes the direction of the price movement so that both positive and negative returns are treated equally, allowing the magnitude of price fluctuations to be captured.

Table 2 lists the variables, their constructions, interpretations, and units used in the empirical analysis.

Table 2. Variable definitions and measurement units

Symbol	Construction	Interpretation in this paper	Unit
		NVIDIA signed daily log return	Percentage points
		Nasdaq-100 signed daily log return	Percentage points
		NVIDIA absolute-return volatility proxy	Percentage points
		Nasdaq-100 absolute-return volatility proxy	Percentage points

Note: The volatility measures are ex post daily proxies. They are not estimates of an unobserved conditional variance.

For example, returns of +3 percent and -3 percent both produce proxy values of 3 percentage points. This measure ignores direction and focuses on movement size. In Fig. 1, twenty-day averages of the absolute returns are displayed to make the daily pattern easier to see. All regressions and forecast tests, however, use the original unsmoothed daily observations.

4.3 Descriptive Patterns

NVIDIA was much more volatile than the Nasdaq-100 in every year. Its average absolute daily return was 2.116 percent in 2023, 2.478 percent in 2024, and 2.178 percent in 2025. The corresponding Nasdaq-100 values were 0.904 percent, 0.863 percent, and 0.953 percent. This difference is expected because an index combines many companies,

so gains in some members can offset losses in others. A single company's price is more exposed to company-specific news.

The standard deviation of NVIDIA returns was highest in 2024 at 3.280 percent. In contrast, Nasdaq-100 return standard deviation was highest in 2025 at 1.478 percent. Fig. 1 shows several periods when both volatility series rose together, but it also shows periods when NVIDIA volatility increased much more than index volatility. The largest Nasdaq-100 absolute return in the sample occurred on April 9, 2025, when the index log return was approximately +11.35 percent and NVIDIA's log return was approximately +17.16 percent. Several other unusually large market movements were also concentrated in early April 2025.




Fig. 1 Short-term volatility during the AI boom, 2023–2025. The twenty-day averages are used only to make the daily series easier to read.

The correlation between signed NVIDIA and Nasdaq-100 daily returns was 0.746. This indicates that the two series frequently moved in the same direction. The correlation between their absolute daily returns was 0.583. This second result is more directly related to the paper's vol-

atility question: large NVIDIA days tended to be large Nasdaq-100 days, although the relationship was far from perfect.

Table 3 reports full-sample descriptive statistics for the signed returns and volatility proxies.

Table 3. Full-sample descriptive statistics (percentage points)

Variable	Mean	SD	Minimum	Median	Maximum
NVIDIA signed return	0.342	3.125	-18.595	0.366	21.809
Nasdaq-100 signed return	0.112	1.265	-6.259	0.151	11.353
NVIDIA absolute-return proxy	2.258	2.186	0.000	1.767	21.809
Nasdaq-100 absolute-return proxy	0.907	0.888	0.002	0.693	11.353

Note: Each row contains 751 observations. SD denotes standard deviation.

Table 4 compares the yearly averages and standard deviations, showing how volatility differed across 2023, 2024, and 2025.

Table 4: Yearly descriptive statistics (percent)

Year	Days	NVDA Mean R	NDX Mean R	NVDA Return SD	NDX Return SD
2023	249	2.116	0.904	2.948	1.140
2024	252	2.478	0.863	3.280	1.150
2025	250	2.178	0.953	3.139	1.478

Note: is the absolute daily log return. SD is standard deviation.

5. Empirical Methodology

5.1 Same-Day Association Model

The first model estimates the relationship between NVIDIA volatility and Nasdaq-100 volatility on the same trading day:

$$V_t^{NDX} = \alpha + \beta V_t^{NVDA} + \varepsilon_t \quad (3)$$

The coefficient β measures the average difference in the Nasdaq-100 absolute-return proxy associated with a one-percentage-point difference in the NVIDIA proxy on the same day. A positive β supports H1.

Ordinary least squares (OLS) were used to estimate the coefficients. The reported R^2 describes in-sample fit: it is the share of sample variation accounted for by the fitted linear relationship. It is not a causal contribution and it does not measure forecast accuracy.

This same-day equation is a relationship model, not a forecast. Both variables are observed on day t . It cannot be used before the day's closing prices are known. It also cannot prove that NVIDIA caused the wider market movement because common news and NVIDIA's index membership can affect the result.

5.2 One-Day-Ahead Forecast Models

The forecasting exercise uses only information available at the end of the previous trading day. The baseline model is:

$$V_t^{NDX} = \alpha + \phi V_{t-1}^{NDX} + \varepsilon_t \quad (4)$$

It asks whether yesterday's Nasdaq-100 volatility helps predict today's volatility. The extended model is:

$$V_t^{NDX} = \alpha + \phi V_{t-1}^{NDX} + \gamma V_{t-1}^{NVDA} + \varepsilon_t \quad (5)$$

The added gamma coefficient asks whether yesterday's NVIDIA volatility contributes information beyond yesterday's index volatility. If gamma is positive and the extended model produces smaller forecast errors, NVIDIA provides additional short-term predictive value.

The models were estimated once using 500 observations from 2023–2024. Their coefficients were then held fixed. For each of the 250 trading days in 2025, the forecast used the actually observed lagged values from the preceding trading day; no 2025 outcome was used to re-estimate the coefficients. A nonnegative floor was specified because an absolute return cannot be negative, although neither model produced a negative 2025 forecast.

5.3 Inference, Diagnostics, and Robustness

Daily financial data can exhibit both nonconstant error variance and short-run serial dependence. Therefore, all coefficient tests use Newey–West HAC standard errors with a maximum lag of five trading days, approximately one trading week [19]. The lag length was fixed before comparing results. The paper reports two-sided tests and 95% confidence intervals; statistical significance is not treated as evidence of economic importance.

Three residual checks are reported. The Breusch–Pagan test evaluates homoskedasticity, the Ljung–Box test evaluates residual serial correlation at lags 5 and 10, and the Durbin–Watson statistic gives a descriptive first-order check. These tests do not make the simple models complete, but they show why HAC inference is preferable to conventional OLS standard errors. Robustness is also evaluated by re-estimating the same-day slope separately for each calendar year and after winsorizing both absolute-return proxies at their 99th percentiles. The winsorized model checks whether a few extreme days drive the result.

5.4 Forecast Evaluation

Forecast accuracy was evaluated against both a historical-mean benchmark and the lagged-index baseline. Mean absolute error (MAE) and root mean squared error (RMSE) are:

$$\text{MAE} = \frac{1}{n} \sum_{t=1}^n |V_t - \hat{V}_t| \quad (6)$$

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{t=1}^n (V_t - \hat{V}_t)^2} \quad (7)$$

MAE gives the average size of an error in percentage points. RMSE gives extra weight to large errors and is especially sensitive to missed volatility spikes. The historical-mean benchmark always forecasts the 2023–2024 mean of . To determine whether the extended and baseline models differ beyond random sampling variation, the paper also tests the mean difference in their absolute and squared losses. The loss-differential standard error is HAC with five lags, following the predictive-accuracy comparison principle of Diebold and Mariano [20]. A positive difference means that the extended model has lower loss. The analysis was carried out with a reproducible script. Prices were sorted by date, merged by common trading day, transformed into log returns, and checked for missing values. The script exports every coefficient table, diagnos-

tic, forecast, and figure used below. The study uses only public market data and contains no personal or confidential information, so it raises no human-subject ethical concerns.

6. Results

6.1 Same-Day Regression

The scatter plot in Fig. 2 shows a clear positive relationship. Many observations are concentrated near the lower-left corner because most trading days had modest movements. As NVIDIA's absolute return increased, large Nasdaq-100 movements became more common. However, the points remain spread out, demonstrating that the NVIDIA proxy alone does not account for every market movement.




Fig. 2: Same-day relationship between NVIDIA and Nasdaq-100 volatility. Each point represents one trading day

The estimated same-day equation was:

(8)

The NVIDIA coefficient was 0.2367 with a HAC standard error of 0.0442. Its t -statistic was 5.36, its p -value was below 0.001, and its 95% confidence interval was [0.150, 0.323]. In sample, a one-percentage-point increase in NVIDIA's absolute daily return was associated with an average 0.237-percentage-point increase in the Nasdaq-100 absolute-return proxy. This is a contemporaneous conditional mean comparison, not an estimate of the causal effect of

NVIDIA on the index.

The model's R^2 was 0.340. Thus, the fitted one-variable relationship accounted for about 34% of the sample variation in Nasdaq-100 absolute returns; it did not establish that NVIDIA generated 34% of market volatility. About 66% remained outside the fitted relationship, consistent with the many companies and information sources affecting the index.

Table 5 reports the same-day regression estimates with HAC standard errors and confidence intervals.

Table 5: Same-day absolute-return proxy regression

Variable	Coefficient	HAC SE	95% CI	-value
Constant	0.3724	0.0816	[0.212, 0.533]	
NVIDIA absolute return	0.2367	0.0442	[0.150, 0.323]	

Observations = 751; ; Newey–West HAC standard errors use five lags.

The intercept was 0.372 percent. It represents the model's expected Nasdaq-100 absolute return when NVIDIA's absolute return is zero. The intercept should not be treated as a precise description of a completely motionless NVIDIA day, because there are few observations exactly at zero. Its main role is to position the regression line.

6.2 Forecast Model Estimates

The baseline forecasting model had little in-sample fit in 2023–2024. The coefficient on the previous day's Nasdaq-100 proxy was 0.0227 (HAC SE = 0.0576,), and the model was 0.0005. In this specification, yesterday's index proxy accounted for almost none of today's absolute-return variation.

When previous-day NVIDIA volatility was added, its coefficient was 0.0386 (HAC SE = 0.0185, , ; 95% CI = [0.002, 0.075]). The training-sample increased from 0.0005 to 0.0098. This provides limited in-sample evidence of a small positive lagged association. However, the extended model still accounted for less than 1% of daily proxy variation, so the coefficient alone does not establish useful forecasting power.

This is an important example of why statistical significance and practical usefulness are not the same. A small coefficient can be statistically significant in a dataset with hundreds of observations, but it may still add very little forecasting accuracy. The out-of-sample test is therefore more important than the training p-value.

Table 6 presents the 2023–2024 training-sample estimates for the baseline and extended one-day-ahead forecasting models.

Table 6: Forecast model estimates, 2023–2024 training sample

Model	Variable	Coefficient	HAC SE	-value
Baseline	Constant	0.8645	0.0575	
Baseline	Previous Nasdaq-100 volatility	0.0227	0.0576	0.694
Extended	Constant	0.8267	0.0606	
Extended	Previous Nasdaq-100 volatility	-0.0349	0.0568	0.540
Extended	Previous NVIDIA volatility	0.0386	0.0185	0.037

Both models use 500 training observations. Baseline ; extended . HAC standard errors use five lags.

6.3 Forecast Performance in 2025

The historical-mean benchmark produced a 2025 MAE of 0.6934 and RMSE of 1.1301 percentage points. The lagged-index baseline was only slightly better, with MAE of 0.6914 and RMSE of 1.1243. The extended model's MAE was 0.6886 and its RMSE was 1.1292. Relative to the baseline, adding NVIDIA reduced MAE by 0.395%

but increased RMSE by 0.443%.

These mixed results do not support a reliable forecasting advantage. The mean baseline-minus-extended absolute-loss difference was only 0.0027 percentage points (), while the squared-loss difference was (). Thus, neither comparison rejected equal predictive accuracy. Fig. 3 illustrates the practical problem: both fitted models remained close to normal daily volatility and missed the large increase during April 2025.




Fig. 3: Actual volatility and one-day-ahead forecasts in 2025. Daily values are averaged within each month for readability.

Table 7 compares out-of-sample forecast errors for 2025 across the historical-mean, baseline, and extended models.

Table 7: Out-of-sample forecast errors for 2025

Model	MAE	RMSE	MAE change vs baseline	RMSE change vs baseline
Historical-mean benchmark	0.6934	1.1301	+0.294%	+0.516%
Baseline: previous Nasdaq-100 volatility	0.6914	1.1243	0.000%	0.000%
Extended: baseline plus previous NVIDIA volatility	0.6886	1.1292	%	+0.443%

Lower errors are better. The 2025 test sample contains 250 trading days. A negative change indicates improvement relative to the lagged-index baseline.

The forecast models appear nearly flat because their lag coefficients are small. This is not a graphing error. It shows that the simple equations mostly predicted a value near the training-period average. They did not have variables describing new information arriving on each day. As a result, they could not anticipate an unusual shock using only yesterday's absolute returns.

6.4 Model Diagnostics and Robustness

Table 8 confirms that conventional OLS inference would be inappropriate for the same-day model. The Breusch–Pagan test strongly rejects homoskedasticity, and the Ljung–Box tests detect residual serial dependence. The forecast-model residuals also show serial dependence at ten lags. HAC standard errors address the resulting uncertainty in coefficient tests, but the diagnostic rejections also indicate that these deliberately simple equations do not capture all time-series dynamics.

Table 8: Residual diagnostics

Model	Durbin–Watson	Ljung–Box p (lag 5)	Ljung–Box p (lag 10)	Breusch–Pagan p
Same-day association	1.848			
Baseline forecast	1.996	0.016	0.001	0.514
Extended forecast	2.006	0.026	0.001	0.523

Note: The Ljung–Box null is no residual serial correlation through the stated lag. The Breusch–Pagan null is homoskedasticity.

The positive same-day slope is not produced only by a few extreme observations. After both proxies were winsorized at their 99th percentiles, the slope was 0.2377,

almost identical to the full-sample estimate of 0.2367. Separate annual estimates were also positive and statistically different from zero in every year. The relationship strengthened from 2023 to 2025, so the pooled estimate should be interpreted as an average over a changing peri-

od rather than as a constant structural parameter.

Table 9 reports robustness checks for the same-day association across the full sample, winsorized sample, and annual subsamples.

Table 9: Robustness of the same-day association

Sample/specification		Slope	HAC SE	95% CI	
Full sample	751	0.2367	0.0442	[0.150, 0.323]	0.340
99% upper-tail winsorized	751	0.2377	0.0178	[0.203, 0.273]	0.334
2023 only	249	0.1509	0.0353	[0.081, 0.220]	0.199
2024 only	252	0.2026	0.0237	[0.156, 0.249]	0.335
2025 only	250	0.3490	0.0989	[0.154, 0.544]	0.487

Note: Winsorization caps each absolute-return proxy at its own 99th percentile. All standard errors are Newey–West HAC with five lags.

Overall, H1 is supported: the same-day slope is positive, precisely estimated with HAC inference, and robust across the reported checks. H2 is not supported. Although the lagged NVIDIA coefficient is positive in the training sample, the 2025 error changes are economically tiny, point in opposite directions under MAE and RMSE, and are not statistically distinguishable from zero in the loss-differential tests.

7. Discussion

The core contribution lies in distinguishing association from forecasting. NVIDIA and the Nasdaq-100 exhibited a strong same-day correlation (0.583), robust under HAC inference, annual subsampling, and winsorization—a result that might misleadingly suggest predictive power. However, 2025 out-of-sample tests and loss-differential p-values confirm this would be an overstatement: lagged NVIDIA volatility yielded no statistically reliable reduction in forecast error.

Two explanations emerge. First, intraday information absorption: investors react to AI-related news within the same trading day, quickly incorporating public information into prices, consistent with market efficiency [15]. Second, mechanical and economic linkage: as a Nasdaq-100 constituent, NVIDIA mechanically drives index returns, while its firm-specific news also influences related semiconductor and AI sectors. This design captures their combined association without disentangling these channels.

Despite inherent limitations, the study's notable strengths: it addresses a salient, intuitive research question; utilizes a high-frequency daily sample; rigorously separates estimation from validation; reports null results honestly; and ensures full reproducibility via the appendix. The practical

implication is clear: importance does not imply predictability. NVIDIA was undeniably central to the AI-boom market narrative, yet its prior-day movement offers no dependable edge for next-day volatility forecasting.

8. Conclusion

This paper examines the relationship between NVIDIA stock-price changes and a Nasdaq-100 absolute-return volatility proxy from 2023–2025. Leveraging 751 daily observations, out-of-sample validation (2023–2024 training, 2025 testing), and Newey–West HAC inference, the study yields a balanced conclusion distinguishing correlation from predictability.

Strong evidence supports H1: NVIDIA and Nasdaq-100 absolute returns correlate at 0.583. Each 1 pp increase in NVIDIA volatility associates with a 0.237 pp rise in market volatility (95% CI = [0.150, 0.323]), explaining 34% of in-sample variation and remaining robust post-winsorization. Conversely, H2 is unsupported; lagged NVIDIA volatility fails to improve next-day forecasts (errors changing <0.5%, MAE/RMSE shifts insignificant). Thus, while tightly coupled intraday—reflecting the AI boom—this relationship offers no dependable predictive edge.

These findings warrant cautious interpretation. First, absolute returns proxy realized volatility but are noisy and distinct from conditional variance metrics (e.g., GARCH). Second, omitting macro variables (rates, earnings, policy) preserves scope but leaves residual dynamics (Ljung–Box significant); HAC standard errors aid inference but preclude causality claims. Third, the 2023–2025 window captures a unique AI cycle; evolving index weights limit generalization. Finally, using unadjusted price indices (NVIDIA post-split continuity notwithstanding) excludes dividend effects. Future work should integrate high-frequency data, macro-financial controls, or event studies (e.g., earnings, AI news) to refine these dynamics.

NVIDIA remains one of the most influential companies in

the AI industry, and its impact on financial markets is likely to continue evolving as AI technologies become more widely adopted. Future studies could examine whether major AI-related events, such as new product launches, earnings announcements, or government AI regulations, strengthen or weaken the relationship between NVIDIA volatility and overall market volatility. In addition, extending the sample period and comparing NVIDIA with other leading AI firms may provide a broader understanding of how the AI sector influences market dynamics over time.

References

- [1] NVIDIA Corporation. Annual Report on Form 10-K for the fiscal year ended January 28, 2024. U.S. Securities and Exchange Commission, 2024. <https://www.sec.gov/Archives/edgar/data/1045810/000104581024000029/nvda-20240128.htm>
- [2] NVIDIA Corporation. Annual Report on Form 10-K for the fiscal year ended January 26, 2025. U.S. Securities and Exchange Commission, 2025. <https://www.sec.gov/Archives/edgar/data/1045810/000104581025000023/nvda-20250126.htm>
- [3] Nasdaq, Inc. Nasdaq-100 Index Methodology. Nasdaq Global Indexes, 2026. https://indexes.nasdaq.com/docs/Methodology_NDX.pdf
- [4] Schwert, G. William. Why Does Stock Market Volatility Change Over Time? *The Journal of Finance*, 1989, 44(5): 1115–1153. <https://doi.org/10.1111/j.1540-6261.1989.tb02647.x>
- [5] Zhang Chao, Yihuang Zhang, Mihai Cucuringu, Zhongmin Qian. Volatility Forecasting with Machine Learning and Intraday Commonality. *Journal of Financial Econometrics*, 2024, 22(2): 492–530. <https://doi.org/10.1093/jffinec/nbad005>
- [6] Zhu Haibin, Lu Bai, Lidan He, Zhi Liu. Forecasting Realized Volatility with Machine Learning: Panel Data Perspective. *Journal of Empirical Finance*, 2023, 73: 251–271. <https://doi.org/10.1016/j.jempfin.2023.07.003>
- [7] Ding Zhuangxin, Clive W. J. Granger, Robert F. Engle. A Long Memory Property of Stock Market Returns and a New Model. *Journal of Empirical Finance*, 1993, 1(1): 83–106. [https://doi.org/10.1016/0927-5398\(93\)90006-D](https://doi.org/10.1016/0927-5398(93)90006-D)
- [8] Engle Robert F. Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica*, 1982, 50(4): 987–1007. <https://doi.org/10.2307/1912773>
- [9] Bollerslev Tim. Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics*, 1986, 31(3): 307–327. [https://doi.org/10.1016/0304-4076\(86\)90063-1](https://doi.org/10.1016/0304-4076(86)90063-1)
- [10] Christensen Kim, Mathias Siggaard, Bezirgen Veliyev. A Machine Learning Approach to Volatility Forecasting. *Journal of Financial Econometrics*, 2023, 21(5): 1680–1727. <https://doi.org/10.1093/jffinec/nbac020>
- [11] Chun Dohyun, Hoon Cho, Doojin Ryu. Volatility Forecasting and Volatility-Timing Strategies: A Machine Learning Approach. *Research in International Business and Finance*, 2025, 75: 102723. <https://doi.org/10.1016/j.ribaf.2024.102723>
- [12] Borochin Paul, Yanhui Zhao. The Economic Value of Equity Implied Volatility Forecasting with Machine Learning. *Journal of Empirical Finance*, 2025, 82: 101618. <https://doi.org/10.1016/j.jempfin.2025.101618>
- [13] Poon Ser-Huang, Clive W. J. Granger. Forecasting Volatility in Financial Markets: A Review. *Journal of Economic Literature*, 2003, 41(2): 478–539. <https://doi.org/10.1257/002205103765762743>
- [14] Nong Huifu. Return and Volatility Connectedness across Stock Markets: A Global Perspective. *Investment Analysts Journal*, 2024, 53(1): 50–71. <https://doi.org/10.1080/10293523.2023.2240562>
- [15] Fama Eugene F. Efficient Capital Markets: A Review of Theory and Empirical Work. *The Journal of Finance*, 1970, 25(2): 383–417. <https://doi.org/10.2307/2325486>
- [16] Zhang Junru, Chen Cui, Chen Zheng, Grantley Taylor. Artificial Intelligence Innovation and Stock Price Crash Risk. *Journal of Financial Research*, 2025, 48(2): 503–543. <https://doi.org/10.1111/jfir.12424>
- [17] Zhou Zhengnan. The Impact of Investor Sentiment Fluctuations on Stock Market Liquidity and Volatility: Implications for Market Returns. *Finance Research Letters*, 2025, 86: 108771. <https://doi.org/10.1016/j.frl.2025.108771>
- [18] nfin. Public financial-data API documentation and historical quote endpoint; upstream provider reported as Nasdaq. Data retrieved June 7, 2026. <https://nfin.dev/>
- [19] Newey Whitney K., Kenneth D. West. A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica*, 1987, 55(3): 703–708. <https://doi.org/10.2307/1913610>
- [20] Diebold Francis X., Roberto S. Mariano. Comparing Predictive Accuracy. *Journal of Business & Economic Statistics*, 1995, 13(3): 253–263. <https://doi.org/10.1080/07350015.1995.10524599>