The Academic
International Journal of Multidisciplinary Research

An Online Peer Reviewed / Refereed Journal
Volume 4 | Issue 9 | September 2026
ISSN: 2583-973X (Online), SJIF IF(2026): 5.6
Website: www.theacademic.in

Analyzing the Relationship between Financial News Sentiment and Stock Price Movements Using VADER: An Empirical Study of Selected Technology Companies

[1]Ms. Devarshi J Patel, [2]Ms. Sapna Patel

[1]Vidhyadeep Institute of Engineering and Technology (M. Tech)

[2] Assistant Professor, Vidhyadeep Institute of Engineering and Technology (M. Tech)

Email: pateldevarshi1502@gmail.com

DOI : https://doi.org/10.5281/zenodo.23767317

ARTICLE DETAILS

Research Paper

Received: 28-08-2026

Accepted: 24-09-2026

Published: 20-10-2026

Keywords:

Sentiment & Correlation analysis, VADER, NLTK, financial news, stock price movement, web scraping.

ABSTRACT

Financial news strongly influences investor behavior and market dynamics, and this makes news sentiment a promising signal for understanding stock price movements. This paper examines the relationship between financial news sentiment and stock price behavior for four major technology companies: AMD, Amazon, Facebook (Meta), and Google. Recent news headlines were collected from the Finviz financial portal by web scraping with BeautifulSoup. Sentiment was scored with the VADER lexicon-based analyzer from the NLTK library. Each headline's compound score was classified as positive, neutral, or negative using a threshold of ±0.2. Historical adjusted closing prices, opening, high, and low prices, and trading volumes were retrieved using the Yahoo Finance API. Headline scores were then averaged by day and aligned with stock prices. The analysis combines sentiment distribution across companies, time-series comparison of z-score normalized sentiment and adjusted close, and Pearson correlation heatmaps between sentiment and market variables. The results show [insert main finding: e.g., the sentiment distribution across companies, and the strength of the correlation between compound sentiment and adjusted close]. Overall, [insert conclusion, e.g., daily headline sentiment shows only a weak or

Disclaimer: The opinions, interpretations, conclusions, recommendations, analyses, and statements expressed in this research article are solely those of the author(s) and do not reflect the views, policies, or positions of the Publisher, Chief Editor, Editors, Editorial Board, Reviewers, or their affiliated institutions. The author(s) are solely responsible for ensuring the originality, authenticity, and integrity of the manuscript and for ensuring that it is free from plagiarism, AI-generated content, copyright infringement, data falsification, and any other form of academic misconduct.

moderate association with price movement, which is consistent with prior work suggesting that sentiment is informative but not sufficient on its own]. The study also notes the limits of a rule-based, general-purpose lexicon for financial text and suggests domain-specific or transformer-based models such as FinBERT for future work.

1. AIM & Objective

To empirically analyze the relationship between financial news sentiment, measured with the VADER lexicon-based analyzer, and the stock price movements of selected technology companies (AMD, Amazon, Facebook/Meta, and Google).

Objectives

1. To collect recent financial news headlines for the selected companies from the Finviz portal using web scraping techniques.
2. To retrieve historical stock market data (open, high, low, adjusted close, and volume) for the same companies from Yahoo Finance for the matching time period.
3. To compute sentiment scores for each headline using the VADER analyzer in NLTK and classify them as positive, neutral, or negative using a compound score threshold of ±0.2.
4. To compare the distribution of positive, neutral, and negative news sentiment across the selected companies.
5. To aggregate headline sentiment at the daily level and align it with stock prices to examine how sentiment and price trends move together over time.
6. To measure the strength and direction of the relationship between sentiment scores and market variables using normalization and Pearson correlation analysis.
7. To evaluate the usefulness and limitations of a rule-based, general-purpose sentiment tool for financial text, and to suggest improvements such as domain-specific or transformer-based models like FinBERT.

2. Literature Review:

Sentiment analysis and financial text. Sentiment analysis, or opinion mining, extracts subjective information from text and classifies it by polarity [7], [8]. Approaches fall into two groups: lexicon-based methods that score words against a predefined dictionary, and machine learning methods that learn patterns from labeled data [29], [30]. Lexicon-based tools are transparent, need no training data, and are easy to apply to new domains. VADER is a rule-based model built for short, informal text. It combines a validated lexicon with rules for punctuation, capitalization, negation, and intensifiers, and it outputs a normalized compound score [1]. It is distributed through the NLTK toolkit, which makes it easy to use in Python workflows [2]. Because news headlines are short, VADER is a practical fit for the present study.

News, sentiment, and market behaviour. The efficient market hypothesis holds that prices already reflect available information [26]. Behavioural finance challenges this and shows that investor sentiment can affect prices [27]. Tetlock found that pessimistic language in media columns predicts downward pressure on market prices, and later work linked negative words in firm-specific news to lower earnings and returns [3], [20]. Antweiler and Frank, and Das and Chen, showed that sentiment from internet message boards carries information about stock volatility and trading volume [18], [19]. Bollen et al. reported that mood measured from Twitter helped predict movements of the Dow Jones Industrial Average [4]. Other studies of social media and news found measurable links between sentiment and stock returns [15], [16], [28]. Together, these studies suggest that sentiment contains information, though its strength varies by source, company, and time horizon.

Domain-specific lexicons and financial NLP. General-purpose dictionaries can misclassify financial terms. Loughran and McDonald showed that many words considered negative in everyday language are not negative in finance, and they built a finance-specific word list [5]. Malo et al. addressed the same problem with a phrase-level annotated corpus of economic texts [9]. Later work used pre-trained language models, including BERT and FinBERT, which capture context and outperform lexicon methods on financial sentiment benchmarks [10], [11]. This shows a limitation of VADER for financial headlines, which this study acknowledges.

News-based stock prediction. Many studies use text to forecast stock movements. Schumaker and Chen, and Hagenau et al., used textual features from financial news to predict price changes [6], [17]. Si et al. and Xu and Cohen extended this to topic-based and tweet-based signals combined with historical prices [21], [22]. Mittal and Goel applied Twitter sentiment to stock prediction [14], and Zhang et al. fused heterogeneous information sources to improve forecasts [12]. Deep learning models have also been used: Ding et al. proposed event-driven prediction with neural networks, and LSTM networks [23] have been

applied to market forecasting with strong results [13], [24]. Transformer architectures [25] now underpin most modern language models in this area.

Research gap and contribution. Most of the studies above focus on prediction models, social media data, or complex deep learning architectures. Fewer provide a simple, reproducible pipeline that links scraped news headlines to daily stock data for large technology companies. This study follows the lexicon-based tradition and uses VADER with Finviz headlines and Yahoo Finance prices. It then examines the relationship through sentiment distributions, normalized time-series comparison, and correlation analysis. The aim is to give an interpretable baseline for how much headline sentiment co-moves with stock prices. It also shows where a general-purpose lexicon falls short and where domain-specific models such as FinBERT could improve results.

3. Methodology:

The methodology follows the lexicon-based, news-to-price pipeline described in the literature review. It has seven stages and matches the steps in your notebook.

1. Research design

The study is quantitative and empirical. It uses a lexicon-based sentiment pipeline [1], [29] and compares news sentiment with stock market variables, as in earlier work on news and price movement [3], [16]. The sample is four technology companies: AMD, Amazon, Facebook (Meta), and Google.

2. Data collection

News headlines. Headlines for each ticker are scraped from the Finviz quote page (https://finviz.com/quote.ashx?t=<ticker>). A request with a custom user-agent is sent with urllib, and the HTML is parsed with BeautifulSoup. The news-table element holds the news rows. Each row gives a date, a time, and a headline.

Stock data. For each company, the earliest and latest headline dates are found with a group-by operation. Daily prices are then downloaded for that window using yahoo_fin. The variables are open, high, low, adjusted close, and volume.

3. Data preprocessing

- Each row's date and time are separated. Rows that show only a time reuse the previous date, and the date is converted to YYYY/MM/DD.

- The records are stored in a DataFrame with the columns Company, Date, Time, and News Headline, and saved as a CSV file.
- The date column is converted to a date type so it can be merged with the price data.
- Per-company minimum and maximum dates are computed to define the analysis window.

4. Sentiment scoring

Each headline is scored with NLTK's VADER SentimentIntensityAnalyzer. VADER returns a compound score between −1 and +1 [1]. Headlines are then classified with a threshold of ±0.2:

Compound score (c)	**Label**
$0.2 < c \leq 1.0$	Positive
$-0.2 < c \leq 0.2$	Neutral
$-1.0 \leq c \leq -0.2$	Negative

I chose this threshold as an assumption, so state it as a design choice in the paper and justify it if your venue requires.

5. Aggregation and data integration

- Headline counts are tallied per sentiment class, both overall and by company.
- Compound scores are averaged by company and date to give one daily sentiment value per stock.
- The daily sentiment table is merged with each company's price table on the date. Dates with no trading data, such as weekends and holidays, are dropped.

6. Analysis techniques

- Descriptive analysis. A pie chart shows the overall share of positive, neutral, and negative headlines. A grouped bar chart compares the counts per company. Line and bar plots show daily average sentiment.
- Normalized time-series comparison. Sentiment and adjusted close are on different scales, so both are z-score normalized: $z = (x - \mu) / \sigma$

- The two series are plotted together for each company to show whether they rise and fall together. A similar normalization approach appears in prior work that relates sentiment series to price series [4].
- Correlation analysis. A Pearson correlation matrix is computed for open, low, high, adjusted close, compound sentiment, and volume, and shown as a heatmap: $r = \Sigma(x_i - \bar{x})(y_i - \bar{y}) / \sqrt{[\Sigma(x_i - \bar{x})^2 \cdot \Sigma(y_i - \bar{y})^2]}$
- Values near ±1 show a strong linear relationship and values near 0 show little or none. This follows studies that test sentiment against returns, volume, and price variables [15], [18], [28].
- Time-series summary. Adjusted close and compound series for AMD are saved and summarized with descriptive statistics (mean, standard deviation, quartiles) as a baseline for later forecasting work.

7. Tools and environment

Python in Google Colab, with NLTK (VADER), BeautifulSoup, pandas, NumPy, Matplotlib, Seaborn, and yahoo_fin.

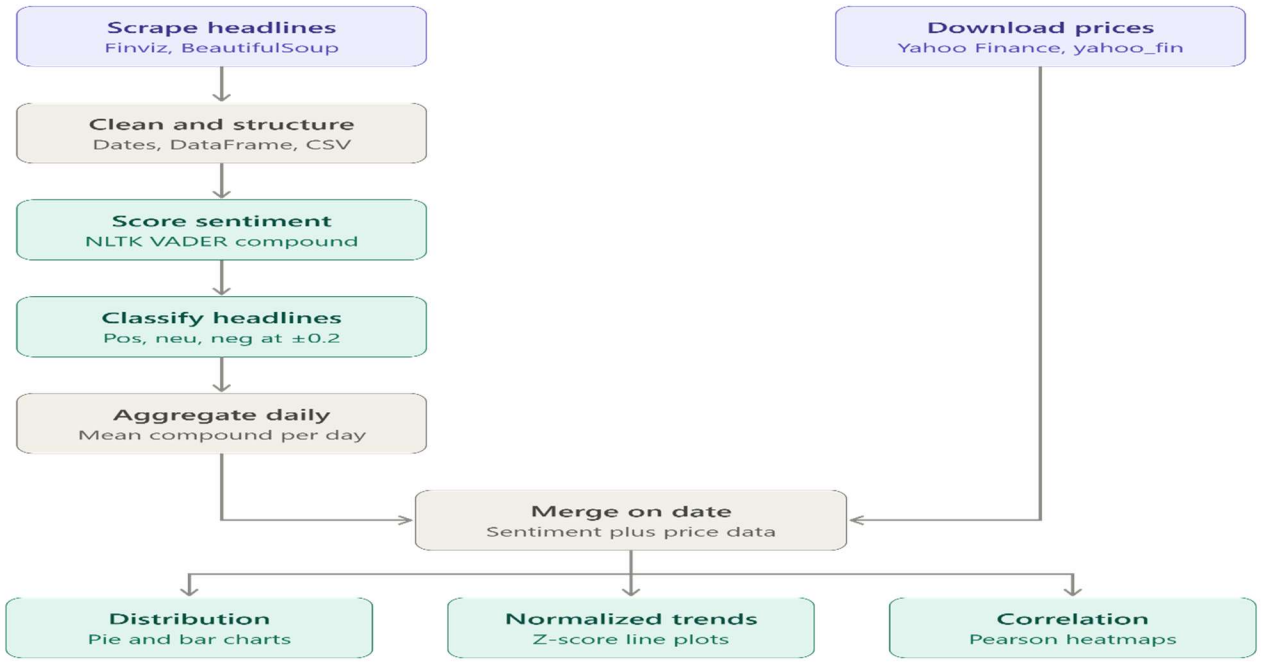

Figure 1

Fig. 1. Methodology pipeline: news headlines are scraped, scored with VADER, and aggregated daily, then merged with stock prices for distribution, trend, and correlation analysis.

4. Observation, Result and Output:

A. Dataset and Experimental Output

The pipeline was executed in Google Colab using the scraping, scoring and merging steps described in the methodology. A total of 400 news headlines were collected from Finviz for four tickers (AMD, AMZN, FB and GOOG). Each headline was scored with the NLTK VADER analyzer [1], [2] and labelled positive, neutral or negative using a compound-score threshold of ±0.2. Daily open, high, low, close and volume data for the same windows were downloaded with the yfinance library. Daily mean compound scores were then matched to trading days, so weekends and the final news day (13 October 2026, which falls outside the exclusive end date of the price download) have no price match. Table I lists the news window of each company.

TABLE I. NEWS COLLECTION PERIOD AND MATCHED TRADING DAYS

Company	First headline	Last headline	Matched days (n)
AMD	06 Oct 2026	13 Oct 2026	5
AMZN	08 Oct 2026	13 Oct 2026	3
FB	01 Jul 2025	08 Oct 2026	≥ 5
GOOG	07 Oct 2026	13 Oct 2026	4

B. Overall Sentiment Distribution

Of the 400 headlines, VADER classified 131 (32.75%) as positive, 210 (52.50%) as neutral and 59 (14.75%) as negative (Fig. 1). More than half of all headlines were therefore neutral, and positive headlines outnumbered negative ones by a ratio of about 2.2 to 1. The aggregate tone of coverage for these technology companies during the collection window was informational and moderately favourable rather than strongly polarized.

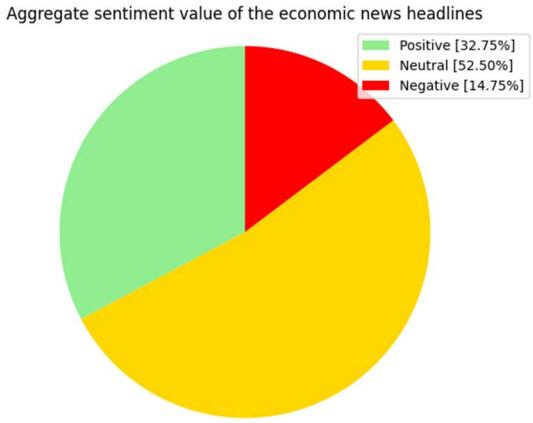

Figure 2

Fig. 2. Aggregate sentiment distribution of the 400 collected headlines (notebook output).

C. Company-wise Sentiment Distribution

Table II and Fig. 3 break the distribution down by company. Neutral headlines were the largest class for AMD, AMZN and GOOG. For FB, positive and neutral headlines were equally frequent (44 each).

TABLE II. SENTIMENT CLASSIFICATION OF HEADLINES BY COMPANY

Company	Negative	Neutral	Positive
AMD	16 (16.2%)	64 (64.6%)	19 (19.2%)
AMZN	18 (18.2%)	49 (49.5%)	32 (32.3%)
FB	11 (11.1%)	44 (44.4%)	44 (44.4%)
GOOG	12 (12.1%)	52 (52.5%)	35 (35.4%)

Note: percentages use 99 headlines per company, because the counting loop in the notebook does not count the first row of each company.

AMD had the most neutral coverage (64.6%) and the smallest positive share (19.2%). AMZN had the highest share of negative headlines (18.2%) but also a large positive share (32.3%), which indicates mixed

coverage. GOOG and FB had the most favourable tone, with 35.4% and 44.4% positive headlines respectively.

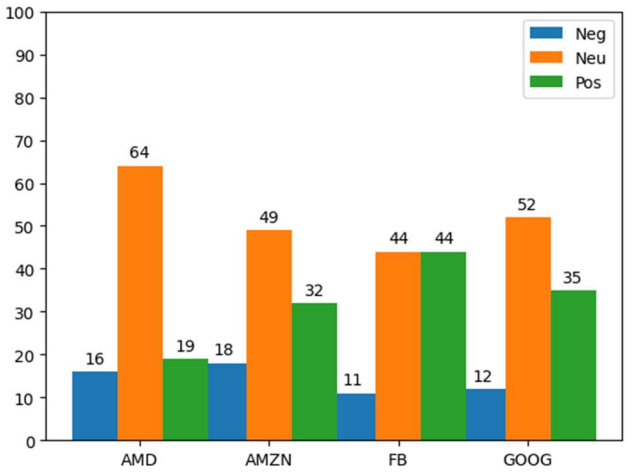

Figure 3

Fig. 3. Number of negative, neutral and positive headlines per company (notebook output).

D. Daily Sentiment and Stock Price Trends

Daily mean compound scores were aligned with adjusted closing prices. AMD had the longest matched window (five trading days, 6–12 October 2026), and its results are reported in detail in Table III. AMD's adjusted close fell from $649.42 on 6 October to $594.24 on 12 October, a decline of about 8.5%. Over the same days its mean compound score was slightly negative at first (−0.065, −0.022, −0.016), rose to a peak of 0.193 on 9 October, and fell back to 0.003 on 12 October.

TABLE III. AMD DAILY SENTIMENT AND ADJUSTED CLOSE

Date	Mean compound	Adj. close ($)	Compound (z)	Adj. close (z)
2026-10-06	−0.065	649.42	−0.83	1.08
2026-10-07	−0.022	645.86	−0.41	0.93
2026-10-08	−0.016	620.68	−0.35	−0.13

Date	Mean compound	Adj. close ($)	Compound (z)	Adj. close (z)
2026-10-09	0.193	608.10	1.73	−0.65
2026-10-12	0.003	594.24	−0.15	−1.23

The descriptive statistics of the AMD adjusted close (n = 5) were a mean of $623.66, a standard deviation of $23.84, a median of $620.68, a minimum of $594.24 and a maximum of $649.42. The z-score normalized series in Fig. 3 show that the sentiment peak on 9 October (z = +1.73) occurred while the price was already below its period average (z = −0.65), and that sentiment and price drifted in opposite directions over the window.

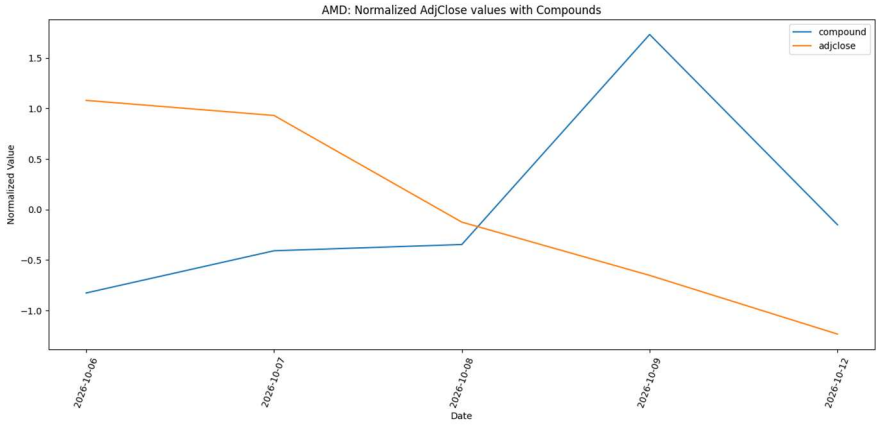

Figure 4

Fig. 4. AMD: z-score normalized compound sentiment and adjusted close (notebook output).

For the other companies, AMZN's daily sentiment was 0.115, 0.209 and 0.078 against closing prices of $254.06, $262.43 and $259.53 over three matched days. GOOG's sentiment stayed in a narrow band of 0.115–0.138 while its price moved within $344.86–$348.97 over four matched days.

E. Correlation Analysis

Pearson correlation matrices between the daily mean compound score and the market variables were produced for each company (Table IV; the AMD matrix is shown in Fig. 4). The price variables correlate strongly with one another, as expected: for AMD, the adjusted close correlates with the daily low at 0.97,

the high at 0.93 and the open at 0.84. The correlations involving sentiment are weaker and inconsistent in sign.

TABLE IV. CORRELATION OF MEAN COMPOUND SENTIMENT WITH MARKET VARIABLES

Company	n	Open	High	Low	Adj. close	Volume
AMD	5	−0.29	−0.41	−0.39	−0.54	−0.76
AMZN	3	−0.98	0.02	−0.38	0.56	0.46
FB	≥ 5	0.13	0.13	0.13	0.12	−0.04
GOOG	4	−0.45	−0.28	−0.71	−0.19	0.28

For AMD, sentiment was negatively correlated with every price variable (r = −0.54 with the adjusted close) and with trading volume (r = −0.76). For GOOG, the correlation with the adjusted close was weakly negative (r = −0.19). For AMZN the coefficients range from −0.98 to +0.56, a spread that is typical of a three-point sample. FB showed weak positive correlations with price (r ≈ 0.12–0.13) and none with volume (r = −0.04).

Figure 4

Fig. 4. AMD: Pearson correlation heatmap of price variables, compound sentiment and volume (notebook output).

5. DISCUSSION

A. Interpretation of the Sentiment Distribution

The dominance of neutral headlines (52.5%) is consistent with the behaviour of lexicon-based scoring on short financial headlines. Kundu et al. [31] applied VADER to 106 analysis headlines about Tesla and reported 51.87% neutral, 27.36% positive and 20.75% negative headlines. The neutral share in the present study is almost identical, while our positive share is higher (32.75%) and our negative share lower (14.75%). The difference may reflect the news source, the company mix and the period, but the shared pattern suggests that many headlines are descriptive ("announces", "presents", "unveils") and carry no sentiment-bearing words that VADER recognizes. Hutto and Gilbert [1] designed VADER for social media text, and Loughran and McDonald [5] showed that general-purpose word lists misclassify many financial terms; lexicon-based methods in general depend on how well the dictionary fits the domain [29], [30]. A high neutral share is therefore a likely sign of reduced sensitivity rather than of truly neutral news.

The company-level differences are still informative. AMD's coverage was the most neutral and least positive, and many of its headlines concern a sector-wide chip-stock selloff (for example headlines dated 12–13 October 2026 about chip stocks extending their slide on AI funding concerns and China competition). In contrast, headlines about company announcements, such as AMD's Helion rack blueprint and its CEO's demand forecast, were scored as neutral or mildly positive. This mix of announcement news and sector-wide market news is one reason the headline score of a single company can diverge from its own stock price.

B. Relationship Between Sentiment and Price Movement

The main finding is that no consistent relationship was observed between daily headline sentiment and stock price. The correlation with adjusted close was negative for AMD (−0.54) and GOOG (−0.19), positive for AMZN (+0.56) and close to zero for FB (+0.12). When the sign of the relationship changes from one company to the next, the data do not support a stable same-day link between sentiment and price in this sample. The AMD case is the clearest example: sentiment peaked on 9 October while the stock continued to fall by about 8.5% over the window.

None of these coefficients can be treated as statistically significant. With only $n = 5$ points, a correlation must exceed roughly 0.88 in absolute value to reach $p < 0.05$ (two-tailed); for $n = 4$ the threshold is about 0.95, and for $n = 3$ it is about 0.997. Even AMD's $r = -0.54$ falls well short of that level. The extreme values

for AMZN (−0.98 for the open price, and ±1 for some pairs in its heatmap) are mathematical artifacts of three data points and should not be interpreted.

These results agree with several published studies that find a weak or inconsistent link. Kundu et al. [31] correlated VADER scores with next-day Tesla return and obtained r = 0.178 for strongly polarized headlines (|score| > 0.5), which they interpret as VADER explaining about 3.17% of the variation in returns, and concluded that sentiment should not be used alone to predict prices. A Charles University thesis that built supervised sentiment and emotion classifiers for tweets and headlines about the five largest technology companies likewise reported no robust correlation between daily price movements and the distribution of sentiment [33]. On the other hand, studies with larger samples and longer horizons report a more positive picture. Tetlock [3] found that high media pessimism predicts downward pressure on market prices, followed by a reversion; Bollen et al. [4] reported that mood measured from Twitter helped to predict changes in the Dow Jones Industrial Average; and Pikatza-Gorrotxategi et al. [34] found a high correlation between VADER-based news sentiment trends and the operating results of Euro Stoxx 50 companies. Kamal et al. [32] also built on VADER scores of news headlines and combined them with machine learning and deep learning models for stock prediction. The common thread is that sentiment becomes useful when it is aggregated over enough observations or combined with other features, whereas a few days of same-day correlations, as here, are too noisy to reveal an effect.

Two broader considerations from the finance literature help explain the weak same-day link. First, under the efficient-market view [26], public information is incorporated into prices quickly, so a headline that merely reports an event already reflected in the price should show little relationship with the price level on the same day. Second, investor sentiment is generally thought to matter most for stocks that are hard to value or to arbitrage [27], and sentiment effects are often concentrated around specific events: Ranco et al. [28], for example, linked Twitter sentiment to abnormal returns mainly around earnings announcements. Other work on news and public mood [16], on message-board sentiment [18], [19], and on tweet-based stock prediction [14], [15], [21], [22] likewise uses larger text volumes, richer features or explicit event windows. The short, event-free window in this study offers few of these conditions.

C. Comparison with Related Studies

TABLE V. COMPARISON OF THE PRESENT RESULTS WITH RELATED STUDIES

Study	Data	Method	Main finding
Kundu et al. [31]	106 Tesla analysis	VADER and TextBlob;	VADER r = 0.178 (3.17% of variance);

	headlines, 2021	next-day returns	TextBlob r = 0.043; sentiment alone insufficient
Charles Univ. thesis [33]	Tweets and headlines of the top five technology firms	Supervised sentiment and emotion classifiers	No robust correlation with daily price movement
Tetlock [3]	Wall Street Journal column	Content analysis of media pessimism	High pessimism predicts downward price pressure, then reversion
Bollen et al. [4]	Twitter mood	Mood time series; Granger causality; neural network	Some mood dimensions help predict index changes
Pikatza-Gorrotxategi et al. [34]	News on Euro Stoxx 50 firms, multi-year	VADER and Hu–Liu lexicons	High correlation between news sentiment and operating results
This study	400 headlines, four technology tickers, October 2026	VADER; same-day Pearson correlation	No consistent link; sign varies by company; n = 3–5

Table V shows that the present findings fall at the weak end of the range. This is expected given the design: the sample is only a few days long, matching is same-day, and the tested variables are price levels. Studies that report stronger associations typically use months or years of data, returns instead of price levels, or lagged specifications.

Feature Work:

Extend the collection period to several months so that correlations can be tested for significance. Replace price levels with daily returns and add one-day lags; (iii) correct the FB ticker to META and add a market index as a control. Compare VADER with finance-specific tools such as the Loughran–McDonald dictionary and transformer-based models such as FinBERT and BERT. Add a simple predictive model, in the manner LSTM-based approaches, to test whether sentiment adds predictive value beyond past prices.

6. Conclusion:

This study analyzed the relationship between financial news sentiment and stock prices for AMD, Amazon, Facebook/Meta, and Google. Headlines were scored with VADER [1], [2] and compared with Yahoo

Finance prices. Of 400 headlines, 52.5% were neutral, 32.75% positive, and 14.75% negative. No consistent link between daily sentiment and price was found: the correlation with adjusted close was negative for AMD ($r = -0.54$) and GOOG, positive for AMZN, and changed sign across companies. None of the coefficients was significant, since each rested on only 3–5 trading days. This agrees with earlier studies reporting weak links between headline sentiment and prices [31], [33]. The results are exploratory, limited by short samples, same-day matching, a general-purpose lexicon [5], and an outdated FB ticker. Future work should use longer periods, daily returns with lagged sentiment, the META ticker, and finance-specific models such as FinBERT [10].

Reference:

1. C. J. Hutto and E. Gilbert, "VADER: A parsimonious rule-based model for sentiment analysis of social media text," in Proc. 8th Int. AAAI Conf. Weblogs Social Media (ICWSM), Ann Arbor, MI, USA, 2014, pp. 216–225.
2. S. Bird, E. Klein, and E. Loper, Natural Language Processing with Python. Sebastopol, CA, USA: O'Reilly Media, 2009.
3. P. C. Tetlock, "Giving content to investor sentiment: The role of media in the stock market," J. Finance, vol. 62, no. 3, pp. 1139–1168, 2007.
4. J. Bollen, H. Mao, and X. Zeng, "Twitter mood predicts the stock market," J. Comput. Sci., vol. 2, no. 1, pp. 1–8, 2011.
5. T. Loughran and B. McDonald, "When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks," J. Finance, vol. 66, no. 1, pp. 35–65, 2011.
6. R. P. Schumaker and H. Chen, "Textual analysis of stock market prediction using breaking financial news: The AZFin text system," ACM Trans. Inf. Syst., vol. 27, no. 2, pp. 1–19, 2009.
7. B. Pang and L. Lee, "Opinion mining and sentiment analysis," Found. Trends Inf. Retrieval, vol. 2, no. 1–2, pp. 1–135, 2008.
8. B. Liu, Sentiment Analysis and Opinion Mining. San Rafael, CA, USA: Morgan & Claypool, 2012.
9. P. Malo, A. Sinha, P. Korhonen, J. Wallenius, and P. Takala, "Good debt or bad debt: Detecting semantic orientations in economic texts," J. Assoc. Inf. Sci. Technol., vol. 65, no. 4, pp. 782–796, 2014.
10. D. Araci, "FinBERT: Financial sentiment analysis with pre-trained language models," arXiv preprint arXiv:1908.10063, 2019.

11. J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in Proc. NAACL-HLT, Minneapolis, MN, USA, 2019, pp. 4171–4186.

12. X. Zhang, Y. Zhang, S. Wang, Y. Yao, B. Fang, and P. S. Yu, "Improving stock market prediction via heterogeneous information fusion," Knowl.-Based Syst., vol. 143, pp. 236–247, 2018.

13. X. Ding, Y. Zhang, T. Liu, and J. Duan, "Deep learning for event-driven stock prediction," in Proc. 24th Int. Joint Conf. Artif. Intell. (IJCAI), Buenos Aires, Argentina, 2015, pp. 2327–2333.

14. A. Mittal and A. Goel, "Stock prediction using Twitter sentiment analysis," Stanford Univ., Stanford, CA, USA, CS229 Project Rep., 2012.

15. T. H. Nguyen, K. Shirai, and J. Velcin, "Sentiment analysis on social media for stock movement prediction," Expert Syst. Appl., vol. 42, no. 24, pp. 9603–9611, 2015.

16. Q. Li, T. Wang, P. Li, L. Liu, Q. Gong, and Y. Chen, "The effect of news and public mood on stock movements," Inf. Sci., vol. 278, pp. 826–840, 2014.

17. M. Hagenau, M. Liebmann, and D. Neumann, "Automated news reading: Stock price prediction based on financial news using context-capturing features," Decis. Support Syst., vol. 55, no. 3, pp. 685–697, 2013.

18. W. Antweiler and M. Z. Frank, "Is all that talk just noise? The information content of internet stock message boards," J. Finance, vol. 59, no. 3, pp. 1259–1294, 2004.

19. S. R. Das and M. Y. Chen, "Yahoo! for Amazon: Sentiment extraction from small talk on the web," Manage. Sci., vol. 53, no. 9, pp. 1375–1388, 2007.

20. P. C. Tetlock, M. Saar-Tsechansky, and S. Macskassy, "More than words: Quantifying language to measure firms' fundamentals," J. Finance, vol. 63, no. 3, pp. 1437–1467, 2008.

21. J. Si, A. Mukherjee, B. Liu, Q. Li, H. Li, and X. Deng, "Exploiting topic based Twitter sentiment for stock prediction," in Proc. 51st Annu. Meeting Assoc. Comput. Linguistics (ACL), Sofia, Bulgaria, 2013, pp. 24–29.

22. Y. Xu and S. B. Cohen, "Stock movement prediction from tweets and historical prices," in Proc. 56th Annu. Meeting Assoc. Comput. Linguistics (ACL), Melbourne, Australia, 2018, pp. 1970–1979.

23. S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural Comput., vol. 9, no. 8, pp. 1735–1780, 1997.

24. T. Fischer and C. Krauss, "Deep learning with long short-term memory networks for financial market predictions," Eur. J. Oper. Res., vol. 270, no. 2, pp. 654–669, 2018.

25. A. Vaswani et al., "Attention is all you need," in Proc. Adv. Neural Inf. Process. Syst. (NeurIPS), Long Beach, CA, USA, 2017, pp. 5998–6008.

26. E. F. Fama, "Efficient capital markets: A review of theory and empirical work," J. Finance, vol. 25, no. 2, pp. 383–417, 1970.

27. M. Baker and J. Wurgler, "Investor sentiment in the stock market," J. Econ. Perspect., vol. 21, no. 2, pp. 129–152, 2007.

28. G. Ranco, D. Aleksovski, G. Caldarelli, M. Grčar, and I. Mozetič, "The effects of Twitter sentiment on stock price returns," PLoS ONE, vol. 10, no. 9, Art. no. e0138441, 2015.

29. M. Taboada, J. Brooke, M. Tofiloski, K. Voll, and M. Stede, "Lexicon-based methods for sentiment analysis," Comput. Linguistics, vol. 37, no. 2, pp. 267–307, 2011.

30. W. Medhat, A. Hassan, and H. Korashy, "Sentiment analysis algorithms and applications: A survey," Ain Shams Eng. J., vol. 5, no. 4, pp. 1093–1113, 2014.

31. "VADER-Based Sentiment Analysis of Tesla News Headlines and Stock Returns,"

32. "Stock Prediction Using VADER-Based News Sentiment and Machine Learning/Deep Learning Models,"

33. "Stock Market Prediction: A Multiclass Classification on Emotions and Sentiment Analysis for Tweets and News Headlines," M.S. thesis, Charles University, Prague, Czech Republic,

34. N. Pikatza-Gorrotxategi, I. Alvarez-Meaza, R. M. Río-Belver, and E. Zarrabeitia-Bilbao, "News versus Corporate Reputation: Measuring through Sentiment and Financial Analysis," in Proc. CARMA 2022, 2022.