



Vaasan yliopisto  
UNIVERSITY OF VAASA

Soumitra Guha

# **How News Headlines Affect Cryptocurrency Prices: A Case Study on Bitcoin and Ethereum**

School of Accounting & Finance

Master's thesis in Finance

Master's Program in Finance

Vaasa 2025

---

**UNIVERSITY OF VAASA****School of Accounting and Finance****Author:** Soumitra Guha**Title of the thesis:** How News Headlines Affect Cryptocurrency Prices: A Case Study on Bitcoin and Ethereum**Degree:** Master of Finance**Discipline:** Master's Degree Programme**Supervisor:** Klaus Grobys**Year:** 2025 **Pages: 93**

---

**ABSTRACT:**

This thesis examines whether curated news headlines are systematically associated with Bitcoin and Ethereum returns in the sample. A mixed-source corpus compiled from CoinDesk and the historical CryptoPanic API is scored with two domain models, FinBERT and CryptoBERT, applied to the duplicate headlines to isolate model effects from data effects. Daily sentiment indices are constructed after cleaning, deduplication, and UTC alignment with prices from Investing.com. Associations are evaluated at three fixed horizons,  $t = 0$ , 7 D, and 30 D, with a limited  $t + 1$  check for forecast evaluation only. Estimation uses ordinary least squares with Newey West heteroskedasticity and autocorrelation consistent standard errors. For weekly and monthly windows, overlap is acknowledged; confirmatory non-overlapping specifications with a simple market factor and realized-volatility control are reported in the appendix.

Results show clear horizon and asset differences. For Bitcoin, FinBERT sentiment is positively associated with same-day returns, while effects fade at 7 D and 30 D; CryptoBERT produces a smaller same-day association. For Ethereum, both models yield positive and statistically significant associations at 7 D and 30 D, with weak or insignificant same-day links. No specification produces a reliable  $t + 1$  effect. The evidence is strictly confirmatory and does not claim causality or tradable prediction.

The contribution is to document, on a single curated headline corpus and under harmonized methods, that model choice and asset differences matter for the timing of sentiment–return associations. The findings inform monitoring and risk governance by identifying when sentiment shocks are most likely to appear in prices.

**Keywords:** Bitcoin, Ethereum, sentiment analysis, FinBERT, CryptoBERT, cryptocurrency returns, market efficiency, behavioral finance

## Table of Contents

1	Introduction	9
1.1	Background of the Study	9
1.2	Problem Statement	10
1.3	Research Objectives	11
1.4	Research Questions	11
1.5	Significance of the Study	11
1.6	Scope and Limitations	12
1.7	Research Gaps	13
1.8	Organization of the Thesis	14
2	Literature Review	15
2.1	Introduction	15
2.1.1	Theoretical foundations for a news–sentiment effect in crypto	15
2.2	Headline Sentiment and Associations with Market Behavior	16
2.3	Thematic Headline Narratives and Economic Context	18
2.4	Headline Networks, Attention, and Spillover Mechanisms	18
2.5	Source Channels, Information Volume, and Cross-Platform Alignment	20
2.6	Forecasting Models, Accuracy, and Methodological Conflicts	20
2.7	Short-Term Predictability and Market Reaction Timing	21
2.8	Regulatory Headlines and Policy Effects	22
2.9	Comparative Sensitivity of Bitcoin and Ethereum	22
2.10	Conflicting Evidence and Contextual Factors	23
2.10.1	Critical comparative analysis and positioning of this thesis	23
2.11	Theoretical Framework	25
2.11.1	Mechanistic link between sentiment and returns	25
2.11.2	Derivation of hypotheses.	26

2.12	Summary of Literature and Hypotheses	28
3	Methodology	29
3.1	Overview	29
3.2	Data Collection	29
3.2.1	Price Data	29
3.2.2	News Headlines and Sentiment Data	29
3.2.2.1	Neutral handling and robustness design	31
3.3	Model Choice: FinBERT and CryptoBERT	31
3.4	Data Preparation	32
3.4.1	Price Data Preparation	32
3.4.2	Return Calculations	32
3.4.3	Sentiment Aggregation	32
3.4.4	Dataset Merging	33
3.5	Econometric Framework	33
3.5.1	Model Specification	33
3.5.2	Estimation Technique	33
3.6	Visualization and Descriptive Analysis	34
3.7	Analytical Workflow	35
3.8	Methodological Justification	35
4	Results and Discussion	36
4.1	Overview	36
4.2	FinBERT Analysis	37
4.2.1	Bitcoin Analysis	37
4.2.1.1	Contemporaneous Relationship Between Sentiment and Returns	39
4.2.1.2	Return Differences by Sentiment Category	41
4.2.1.3	Dynamic Co-movement of Returns and Sentiment	42

4.2.1.4	Limited forecast evaluation	43
4.2.1.5	Discussion of Findings	43
4.2.1.6	Summary of Empirical Evidence	44
4.2.1.7	Concluding interpretation (FinBERT, Bitcoin)	45
4.3	Ethereum Analysis	45
4.3.1.1	Descriptive Analysis of FinBERT Sentiment	46
4.3.1.2	Contemporaneous Relationship Between Sentiment and Returns	47
4.3.1.3	Return differences by sentiment category (FinBERT, Ethereum)	50
4.3.1.4	Dynamic co-movement of returns and sentiment (FinBERT, Ethereum)	51
4.3.1.5	Limited forecast evaluation (FinBERT, Ethereum)	52
4.3.1.6	Discussion of findings (FinBERT, Ethereum)	52
4.3.1.7	Summary of empirical evidence (FinBERT, Ethereum)	53
4.3.1.8	Concluding Interpretation	53
4.4	CryptoBERT Results	54
4.4.1	Bitcoin analysis (CryptoBERT)	54
4.4.1.1	Descriptive analysis of CryptoBERT sentiment (BTC)	54
4.4.1.2	Contemporaneous association between sentiment and returns (CryptoBERT, Bitcoin)	57
4.4.1.3	Return differences by sentiment category (CryptoBERT, Bitcoin)	60
4.4.1.4	Dynamic co-movement of returns and sentiment (CryptoBERT, Bitcoin)	61
4.4.1.5	Limited forecast evaluation	62
4.4.1.6	Discussion of Findings	62
4.4.1.7	Summary of Empirical Evidence	63
4.4.1.8	Concluding interpretation (CryptoBERT, Bitcoin)	63
4.5	Ethereum analysis (CryptoBERT)	64
4.6	Cross-Model and Cross-Asset Comparison	71

4.7	Robustness and Consolidated Results	73
4.8	Including neutral headlines: robustness check	74
5	Conclusion	76
	References	80
	Appendix	90

## List of Tables

Table 1: descriptive statistics of the FinBERT sentiment variable .....	37
Table 2: Regression Results (Contemporaneous and Predictive Models).....	39
Table 3: Descriptive Statistics of FinBERT Sentiment (ETH).....	46
Table 4: Regression Results - FinBERT Sentiment and ETH Returns.....	48
Table 5: Descriptive Statistics of CryptoBERT Sentiment (BTC).....	55
Table 6: Regression Results CryptoBERT Sentiment and BTC Returns .....	57
Table 7: Summary of Empirical Evidence.....	63
Table 8: Ethereum Analysis.....	64
Table 9: Consolidated results across models, assets, and horizons .....	73

## List of Figures

Figure 1: Mechanism linking headline sentiment to cryptocurrency returns. ....	27
Figure 2: Distribution of FinBERT Sentiment .....	37
Figure 3: BTC Daily Return vs FinBERT Sentiment .....	39
Figure 4: BTC Daily Returns by Sentiment Category.....	41
Figure 5: BTC Returns and Sentiment Over Time .....	42
Figure 6: Distribution of FinBERT Sentiment (ETH).....	47
Figure 7: ETH Daily Return vs FinBERT Sentiment .....	48
Figure 8: ETH Daily Returns by FinBERT Sentiment Category.....	50
Figure 9: ETH Returns and FinBERT Sentiment Over Time .....	51
Figure 10: Distribution of CryptoBERT Sentiment (BTC).....	55
Figure 11: BTC Daily Return vs CryptoBERT Sentiment .....	58
Figure 12: BTC Daily Returns by CryptoBERT Sentiment Bucket .....	60
Figure 13: BTC Returns and CryptoBERT Sentiment Over Time .....	61
Figure 14: Distribution of CryptoBERT Sentiment .....	66
Figure 15: ETH Daily Returns by Sentiment Bucket .....	67
Figure 16: ETH Daily Return vs. Sentiment Scatter .....	68
Figure 17: ETH Returns and Sentiment Over Time .....	69

# 1 Introduction

## 1.1 Background of the Study

Digital markets now transmit information at high speed, and prices often respond to the arrival of public news and the sentiment it conveys. In this environment, the growth of online news outlets and algorithmic trading has reshaped how market participants process headlines and incorporate them into orders. Cryptocurrencies have been a prominent part of this shift. They introduced new forms of value representation and new channels through which information can affect prices.

Bitcoin was introduced in 2009 and Ethereum followed in 2015 with programmable smart contracts. These assets trade continuously across global venues, and they lack the conventional anchors that guide valuation in many traditional markets. As a result, they are more exposed to attention dynamics, liquidity constraints, and limits to arbitrage. High volatility attracts both institutional and retail participation, which can amplify short run responses to news.

Headline content is a salient input for attention formation and order flow. Curated news headlines provide time stamped signals with clearer editorial standards relative to social feeds, and they reach broad audiences quickly. Short headlines can coincide with abrupt changes in order imbalance and volatility. Prior work also indicates that tone may map to reactions that are asymmetric, with negative news sometimes producing larger absolute effects than positive news. Since Bitcoin and Ethereum are closely followed and often co owned, common headlines and market wide narratives can appear as cross asset comovement.

The present study uses curated daily headlines to examine how headline tone and related thematic signals are associated with Bitcoin and Ethereum market outcomes. We adopt a confirmatory perspective. Our goal is to document associations at clearly defined horizons rather than to claim causal effects or tradable alpha. Throughout the thesis we maintain consistent definitions of variables and windows. Sentiment is measured as continuous daily indices constructed from headlines. Market outcomes are evaluated at same day, seven day, and thirty day horizons, with a limited lead lag check at  $t + 1$  for forecast evaluation.

Estimation relies on ordinary least squares with Newey West heteroskedasticity and autocorrelation consistent standard errors. Where overlapping windows are relevant, we acknowledge the overlap and focus on robust inference. This framing aligns the study with attention based mechanisms and the reality of 24 by 7 crypto trading while keeping all claims strictly confirmatory.

## **1.2 Problem Statement**

Existing research on text sentiment and asset prices is extensive, yet the specific role of curated news headlines in cryptocurrency markets remains insufficiently understood. Headlines are the first contact point for many investors and shape attention and expectations before the underlying articles are read. In a 24 by 7 market with high volatility and heterogeneous participants, headline signals can coincide with sharp changes in order flow and price.

Most prior studies aggregate broader news or rely on social media streams. These choices introduce timestamp noise, variable editorial quality, and mixture of topics that can obscure the unique effect of headline language. Findings also diverge on directionality. Some studies report that sentiment helps explain or forecast prices, while others document that price movements feed back into measured sentiment. The variation in data sources, horizons, and estimation choices contributes to the lack of consensus.

This study addresses that gap with a confirmatory design. The objective is to assess the in sample association between daily headline sentiment and Bitcoin and Ethereum market outcomes under clearly defined horizons. We examine same day, seven day, and thirty day windows, and we include a limited  $t + 1$  check for forecast evaluation without making claims about tradable alpha or causality. Estimation uses ordinary least squares with Newey West heteroskedasticity and autocorrelation consistent standard errors. Overlap in multi day windows is acknowledged and handled through robust inference. By focusing on curated headlines and applying two sentiment classifiers on the same corpus, the study provides a disciplined comparison that clarifies whether headline tone is systematically associated with returns and volatility for Bitcoin and Ethereum.

### 1.3 Research Objectives

- a. Examine the in-sample association between continuous daily headline sentiment (FinBERT and CryptoBERT) and Bitcoin and Ethereum returns at  $t = 0$ , 7D, and 30D.
- b. Conduct a limited  $t + 1$  lead lag check as a forecast evaluation exercise, without claims of predictive power or tradable alpha.
- c. Compare association strength across BTC and ETH, and across FinBERT and CryptoBERT, under harmonized controls and identical specifications.
- d. Maintain a confirmatory scope with clear definitions, HAC-robust inference, and transparent handling of overlapping multi day windows.

### 1.4 Research Questions

To achieve these objectives, the study is guided by the following research questions:

1. How is continuous daily headline sentiment associated with same day ( $t = 0$ ) and cumulative (7D, 30D) returns of BTC and ETH in sample?
2. Does today's sentiment exhibit any  $t + 1$  lead lag association with next day returns in sample, recognizing that this is a limited forecast evaluation only?
3. How do the association magnitudes and directions differ between BTC and ETH and between FinBERT and CryptoBERT when modeling choices are held constant?
4. Are conclusions consistent across the two classifiers applied to the same headline corpus, and when overlapping multi day windows are addressed with HAC-robust inference?

### 1.5 Significance of the Study

This study contributes to behavioral finance and crypto economics by isolating the role of curated news headlines in shaping price formation for Bitcoin and Ethereum. Unlike many traditional assets, major cryptocurrencies trade 24 by 7 and lack conventional cash flow anchors, which heightens sensitivity to attention, liquidity conditions, and limits to arbitrage. By focusing on headline text and clearly defined return horizons, the study documents associations that clarify how short text sentiment relates to market outcomes in a confirmatory setting.

The findings carry practical value for traders and risk managers. Headline tone can inform monitoring, risk limits, and post news positioning when used alongside existing controls and with appropriate caution. For policymakers and regulators, the results illustrate how public communication and policy signals can coincide with measurable changes in widely held digital assets. For researchers, the work narrows the gap between sentiment analysis and market microstructure by using a transparent design. The use of a curated headline corpus, two sentiment classifiers applied to the same data, ordinary least squares with Newey West standard errors, and explicit treatment of overlapping windows supports reproducibility and disciplined comparison without causal claims.

## **1.6 Scope and Limitations**

This study focuses on Bitcoin and Ethereum. These two assets are the most liquid and widely covered by professional news outlets, which supports reliable headline collection and consistent measurement. The analysis uses curated online financial news headlines as the sole text source. Social media content and long form articles are not included.

The work concentrates on text. Non textual media such as images, videos, and memes are outside scope. Sentiment is measured as continuous daily indices derived from headlines using FinBERT and CryptoBERT applied to the same corpus. Outcomes are limited to returns at  $t = 0, 7D, \text{ and } 30D$ , with a limited  $t + 1$  check for forecast evaluation. The study does not claim causality or tradable alpha.

Time alignment uses daily aggregation with UTC based matching of headline timestamps to return windows. In continuously trading markets, publication time and market reaction may not be perfectly synchronized at the intraday level. This timing uncertainty is a limitation that is mitigated by focusing on daily windows and by reporting robust inference.

Estimation relies on ordinary least squares with Newey West heteroskedasticity and autocorrelation consistent standard errors. Multi day windows can overlap, which increases serial correlation. This is addressed with HAC standard errors, and results are interpreted as confirmatory associations within sample.

Findings are specific to Bitcoin and Ethereum under the defined period and data source. They should not be generalized to smaller or less liquid cryptocurrencies without further evidence. The study does not evaluate trading strategies, transaction costs, or execution feasibility.

## 1.7 Research Gaps

Existing crypto-sentiment studies commonly rely on social media streams rather than curated news headlines, evaluate a single sentiment model in isolation, and focus on one asset and one horizon. As a result, it is unclear whether finance oriented and crypto oriented language models extract meaningfully different signals from the same headline corpus, and whether any such differences matter for in sample associations with Bitcoin and Ethereum returns (Gurgul et al., 2023; Kirtac & Germano, 2024; Loginova et al., 2021; Moradi-Kamali et al., 2025; Nie et al., 2024).

This thesis addresses that gap within a confirmatory scope.

1. It constructs a unified headline dataset for Bitcoin and Ethereum, cleaned, deduplicated, and time aligned to market data so that any differences in results can be attributed to the model rather than to differing inputs.
2. It applies two transformer classifiers, FinBERT and CryptoBERT, to the same corpus with consistent handling of polarity classes and daily aggregation to form continuous sentiment indices.
3. It links headline sentiment to returns at  $t = 0$ , 7D, and 30D using ordinary least squares with Newey West heteroskedasticity and autocorrelation consistent standard errors. Overlapping multi day windows are acknowledged, and inference is kept robust. A limited  $t$  plus 1 check is included for forecast evaluation without claims of predictive power or tradable alpha.
4. It runs all specifications in parallel for Bitcoin and Ethereum under harmonized modeling choices to reveal asset specific differences in association strength and horizon.

By holding the information source constant while varying the classifier and the asset, the thesis provides a comparative, model sensitive answer to a question left open by prior work. When headlines are the common input, do FinBERT and CryptoBERT yield economically and statistically meaningful differences in their in-sample association with Bitcoin and Ethereum

returns, and on which horizons. The explicit mapping from prior limitations to this design constitutes the study's contribution, framed strictly as confirmatory evidence.

## **1.8 Organization of the Thesis**

This thesis is structured in five chapters.

Chapter 1 introduces the study. It states the background, problem statement, objectives, research questions, significance, scope, limitations, research gaps, and the organization of the document.

Chapter 2 presents the literature review and theoretical framework. It analyzes prior work on sentiment and cryptocurrency markets, explains why results differ across sources and methods, and outlines the core mechanisms relevant to this study, including attention, limits to arbitrage, and continuous trading. The chapter positions curated headlines as the information source and motivates a confirmatory approach.

Chapter 3 details the methodology and data. It describes price and headline sources, corpus construction, preprocessing, FinBERT and CryptoBERT application, daily sentiment aggregation, and UTC based alignment with return windows. It then sets out the econometric design using ordinary least squares with Newey West heteroskedasticity and autocorrelation consistent standard errors, notes overlap in multi day windows, and explains the specification choices used for all comparisons. Visual summaries are included for transparency and reproducibility.

Chapter 4 reports the empirical results. It presents associations between headline sentiment and returns at  $t = 0, 7D, \text{ and } 30D$ , a limited  $t + 1$  check for forecast evaluation, and comparative evidence across classifiers and across assets. The chapter discusses robustness notes consistent with the confirmatory scope and interprets magnitudes in light of the theoretical mechanisms.

Chapter 5 concludes. It summarizes the main findings, discusses theoretical and practical implications, states limitations, and identifies directions for future research.

## 2 Literature Review

### 2.1 Introduction

Information arrives continuously in cryptocurrency markets and is rapidly incorporated into trading decisions. Headlines are a prominent channel because they are brief, time stamped, and widely disseminated. Bitcoin and Ethereum attract persistent coverage, so headline tone can coincide with short run changes in attention, order flow, and prices. Prior work reports effects on returns and volatility, but the direction, magnitude, and persistence of these effects vary with the data source, labeling approach, and estimation choices (Ahmad et al., 2015; Anese et al., 2023; Chen et al., 2021; Guidolin & Pedio, 2021; Heston & Sinha, 2016; Liu et al., 2022).

Differences across studies often reflect three factors. First, the information source varies. Social media streams introduce timestamp noise and heterogeneous editorial quality, while curated news headlines provide clearer timing and scope. Second, model and feature choices differ. Studies use distinct sentiment pipelines, polarity definitions, and aggregation rules, which can change measured tone. Third, design choices matter. Horizons, overlapping windows, and treatment of serial correlation influence inference. These elements explain why some studies find that sentiment helps explain prices while others find that price movements feed back into measured sentiment (Baker & Wurgler, 2006; Huang & Ibragimov, 2022; Mai et al., 2022; Sakariyahu et al., 2023; Symitsi & Stamolampros, 2021).

This chapter reviews the literature with that structure in mind. It contrasts social media and curated news evidence, evaluates findings across models and horizons, and links results to mechanisms grounded in attention, limits to arbitrage, and continuous trading. The goal is to clarify how methodological choices shape conclusions and to motivate a confirmatory design that focuses on association, uses curated headlines as the common input, and applies consistent specifications across Bitcoin and Ethereum.

#### 2.1.1 Theoretical foundations for a news–sentiment effect in crypto

Crypto markets have features that make headlines influential as signals for short horizon price formation.

First, continuous trading removes overnight closures and batch openings that can delay price discovery in other markets. With trading active at all hours, order flow can incorporate headline tone on the same day. This motivates testing  $t = 0$  effects using daily headline aggregates and a limited  $t + 1$  check for forecast evaluation (Aït-Sahalia et al., 2024; Beschwitz et al., 2019; Brière et al., 2022; Chen et al., 2021; Deveikyte et al., 2022).

Second, limits to arbitrage are often stronger in crypto. Liquidity is fragmented across venues, funding conditions vary across stablecoins and derivatives, and occasional outages or frictions restrict capital mobility. When arbitrage capacity is constrained, sentiment driven mispricing can persist long enough to be observable at daily horizons. This justifies the use of ordinary least squares with Newey West heteroskedasticity and autocorrelation consistent standard errors to handle serial correlation from gradual adjustment (Kommel et al., 2018; Szczygielski et al., 2020).

Third, attention theory predicts that investors respond more to salient, easy to process cues. Headlines are short, prominent, and arrive in bursts. On high attention days, coordination of focus can amplify price impact. This is why the empirical design treats curated headlines as the information source and includes headline counts in descriptive summaries and, where applicable, as a control (Aït-Sahalia et al., 2024; Beschwitz et al., 2018; Dim et al., 2023).

Fourth, investor clientele and narrative intensity can generate asymmetric reactions. Technology adoption, regulation, and liquidity stories may elicit stronger responses to certain tones or assets. This provides an interpretive basis for sign and magnitude differences between Bitcoin and Ethereum and for evaluating whether effects attenuate from  $t = 0$  to 7D and 30D (Augustin et al., 2023; Meegan et al., 2020; Ortu et al., 2021).

These mechanisms guide the confirmatory design. The study tests clearly defined associations between daily headline sentiment and returns, uses curated headlines for cleaner timing, and applies consistent specifications across assets without making causal or alpha claims (Baker & Wurgler, 2006; Lefort et al., 2024).

## **2.2 Headline Sentiment and Associations with Market Behavior**

A large strand of research reports that adding sentiment features to price based models improves forecasts for Bitcoin and Ethereum. Studies using neural networks such as MLP,

CNN, and LSTM often find lower forecast errors once sentiment inputs are included, which suggests that text captures behavioral dimensions not fully reflected in lagged prices. Similar results appear during turbulent periods. For example, work that combines social media sentiment and search intensity with time series models such as SARIMA reports improved fit for BTC and ETH, and transformer based classifiers can outperform earlier sentiment tools in those forecasting settings (Arslan, 2024; Chalkiadakis et al., 2023; Ciganovic & D'Amario, 2023; Gurgul et al., 2023; Hossain et al., 2024; Liapis et al., 2021; Moradi-Kamali et al., 2025).

Other frameworks separate the prediction and validation steps. In these designs, one model forecasts prices and a second component checks whether sentiment supports the predicted direction. Such hybrids can reduce forecast errors relative to price only baselines (Chen et al., 2024).

At the same time, several papers show that the incremental value of sentiment weakens once broader market factors are controlled or when sentiment is measured from sources that may respond to price rather than lead it. When overall market movement is included, measured lead lag links between sentiment and next day returns can shrink, which implies that part of the signal reflects contemporaneous mood rather than independent information (Allen et al., 2019; Deveikyte et al., 2022; Gan et al., 2019; Mai et al., 2022).

Taken together, the literature shows promise for sentiment features in forecasting exercises but also highlights sensitivity to data source, horizon, and controls. Social media streams and mixed indicators can raise timestamp noise and introduce feedback from price to measured sentiment. This motivates the present thesis to adopt a confirmatory design based on curated daily headlines, to focus on clearly defined in sample associations with returns at  $t = 0, 7D,$  and  $30D,$  to include only a limited  $t + 1$  check for forecast evaluation, and to apply consistent specifications across Bitcoin and Ethereum. This approach clarifies whether headline tone from a clean source is systematically associated with returns without making causal or trading claims (Bashchenko, 2022; Gadi & Sicilia, 2024; Moradi-Kamali et al., 2025; Said et al., 2023).

### **2.3 Thematic Headline Narratives and Economic Context**

Headline topics shape how information is processed in crypto markets. Prior studies classify cryptocurrency news into recurring themes such as regulation, security incidents, governance, and investment. Negative items within these themes, for example enforcement actions or hacking events, are consistently associated with short run price declines, which underscores the salience of negative framing in attention and order flow (Akyildirim et al., 2024; Chokor & Alfieri, 2021; Coulter, 2022; Muktadir-Al-Mukit & Ali, 2025; Zhang et al., 2025).

Narrative clustering typically highlights investment, technology, regulation, and security as major groups. Reported relationships with prices differ by theme and by asset. Some studies find stronger links for investment and regulation narratives, with weaker or asymmetric links for technology and security. These patterns suggest that the market reacts more to headlines that change perceived participation, access, or legal risk than to technical updates that are slower to value (Davoudi et al., 2024; Jesus & Dumitrescu, 2025; Meegan et al., 2020; Schwenkler & Zheng, 2025).

Broader macroeconomic news also conditions crypto returns through the sentiment embedded in financial headlines. During the COVID 19 period, several papers document that indices such as Economic News Sentiment and the VIX are associated with movements in Bitcoin and Ethereum, with differences in sensitivity across the two assets. These findings indicate that headline tone can combine asset specific narratives with macro risk and optimism signals (Canayaz et al., 2023; Filippou et al., 2023; Meegan et al., 2020).

This thesis does not model narrative themes separately. The literature in this section is used to interpret results and to motivate the focus on curated headlines and clearly defined horizons. The empirical design remains confirmatory. It tests whether continuous daily headline sentiment is systematically associated with returns for Bitcoin and Ethereum, without causal claims or trading assertions.

### **2.4 Headline Networks, Attention, and Spillover Mechanisms**

Beyond direct tone effects, headline linkages can transmit shocks across assets. Bitcoin and Ethereum occupy central positions in crypto news ecosystems and appear frequently as co

mentions. When one asset experiences a negative headline shock, co mentioned peers can exhibit short run movements that later revert. This pattern is consistent with temporary mispricing driven by correlated attention rather than by shared fundamentals (Business-Level Strategy, 2020; Ge et al., 2024; Han et al., 2022; Schwenkler & Zheng, 2025).

Attention intensity amplifies these dynamics. Weeks with spikes in Bitcoin and Ethereum coverage are often accompanied by surges in peer trading activity and return co-movement, which indicates that concentrated media focus can synchronize behavior across assets. In this sense, Bitcoin and Ethereum function as informational anchors within the headline network, and their coverage can shape how other cryptocurrencies are valued in the short run (Corbet et al., 2018; Meegan et al., 2020).

This thesis does not model network spillovers directly. The network evidence informs interpretation only. Our empirical design remains confirmatory and asset specific. We test whether continuous daily headline sentiment is associated with Bitcoin and Ethereum returns at clearly defined horizons, and we avoid causal or trading claims. Beyond direct tone effects, headline linkages can transmit shocks across assets. Bitcoin and Ethereum occupy central positions in crypto news ecosystems and appear frequently as co-mentions. When one asset experiences a negative headline shock, co mentioned peers can exhibit short run movements that later revert. This pattern is consistent with temporary mispricing driven by correlated attention rather than by shared fundamentals.

Attention intensity amplifies these dynamics. Weeks with spikes in Bitcoin and Ethereum coverage are often accompanied by surges in peer trading activity and return co-movement, which indicates that concentrated media focus can synchronize behavior across assets. In this sense, Bitcoin and Ethereum function as informational anchors within the headline network, and their coverage can shape how other cryptocurrencies are valued in the short run (Corbet et al., 2018).

This thesis does not model network spillovers directly. The network evidence informs interpretation only. Our empirical design remains confirmatory and asset specific. We test whether continuous daily headline sentiment is associated with Bitcoin and Ethereum returns at clearly defined horizons, and we avoid causal or trading claims.

## 2.5 Source Channels, Information Volume, and Cross-Platform Alignment

Headline effects depend on the source. Social media carries large volumes of sentiment but often exhibits higher noise and faster feedback from price to measured tone. Institutional news sources provide clearer timing, editorial standards, and topic focus, which can yield more persistent market responses. This contrast motivates the thesis focus on curated headlines as the primary information channel (Brière et al., 2022; Bybee et al., 2024).

Cross platform agreement can strengthen signals. Studies report that when traditional media and social media convey similar tone, short horizon associations with returns become more pronounced. By contrast, sentiment drawn from a single social stream has lower explanatory power in many settings. Engagement metrics such as likes and retweets correlate with price movements, yet realized profitability varies with asset and conditions, which highlights the risk of relying on platform activity alone (Kant et al., 2024; Sundarasan & Saleem, 2025).

Information volume matters alongside polarity. Neutral or information dense headlines can affect liquidity and trading even without strong positive or negative tone, which suggests that investors respond to attention and content load as well as to sentiment direction. In this thesis, the empirical design remains headline based and confirmatory. We use curated daily headlines to construct continuous sentiment indices, maintain clearly defined horizons, and treat headline counts as descriptive context or, where applicable, as a simple control. The goal is to isolate associations from a clean source without making causal or trading claims (Binsbergen et al., 2024; Brière et al., 2022).

## 2.6 Forecasting Models, Accuracy, and Methodological Conflicts

Headline based forecasting has advanced with deep learning and modern NLP. Studies applying transformer models such as FinBERT and sequence models such as Bi-LSTM report higher classification and forecasting accuracy than lexicon approaches in several datasets. Weighted sentiment schemes that scale tone by source credibility or audience reach can further improve forecast metrics in model specific settings (Karzanov, 2023; Liapis et al., 2021; Zhu et al., 2022).

Comparative studies often find that incorporating sentiment features raises directional accuracy or Sharpe ratios relative to price only baselines. These gains, however, depend on data scope, feature engineering, and evaluation design. Reported improvements are typically measured against in sample or rolling benchmarks and may be sensitive to class balance, timestamp alignment, and how overlapping horizons are handled (Chuang et al., 2024; Kisiel & Gorse, 2022; Li et al., 2021).

Disagreement persists on causality and directionality. Some nonlinear frameworks, including Gaussian process models, detect two way links between sentiment and returns. Linear VAR and related entropy based analyses often find the opposite direction, with prices leading measured sentiment. These conflicting results can arise from differences in information sources, horizon choices, and model flexibility. Nonlinear methods can capture feedback and delayed effects, while static or linear approaches can understate them, especially when inputs contain noise or when windows overlap (Calderon & Berman, 2024; Ghorbani et al., 2024; Gonçalves et al., 2024).

## **2.7 Short-Term Predictability and Market Reaction Timing**

Short horizon reactions to headline tone are frequently the strongest in the literature. Studies document that negative headlines are often followed by immediate drawdowns, while positive headlines can coincide with quick upticks in prices and volatility. Classification frameworks that flag impact based on intraday price movements also report high directional accuracy when using modern transformer representations of headline text. At the same time, several papers find temporal asymmetry. Good news can diffuse more slowly across investors, while negative news tends to trigger faster and larger immediate responses. This asymmetry helps explain why short run effects can differ by sign and why persistence can vary across horizons (Baker & Wurgler, 2006; Calvet & Fisher, 2008; Hirshleifer et al., 2023; Tao & Shao, 2025).

The present thesis treats these findings as context for interpretation. Our empirical design is confirmatory and focuses on clearly defined daily windows. We test associations between daily headline sentiment and returns at  $t = 0, 7D,$  and  $30D,$  and include a limited  $t + 1$  check for forecast evaluation. We do not claim causality or trading viability. Estimation uses ordinary

least squares with Newey West heteroskedasticity and autocorrelation consistent standard errors, and overlap in multi day windows is acknowledged.

## **2.8 Regulatory Headlines and Policy Effects**

Regulatory communication is a central narrative channel for cryptocurrency markets. Headlines that indicate supportive policies, institutional adoption, or legal clarity are commonly associated with positive short run market reactions. Restrictive or uncertain policy signals are associated with declines. Prior work also shows that policy news can coincide with changes in activity measures on public blockchains, which suggests that regulatory communication affects both valuation and participation (Chokor & Alfieri, 2021; Meegan et al., 2020; Shanaev et al., 2019).

Sensitivity to policy appears broad based across significant assets, including Bitcoin and Ethereum. This supports the view that regulatory signals operate at the ecosystem level rather than at the level of individual tokens alone. In this thesis, regulatory headlines are part of the curated corpus but are not modeled as a separate category. The regulatory literature informs interpretation only. Our analysis remains confirmatory. We test whether continuous daily headline sentiment, constructed from all curated headlines, is associated with returns at the specified horizons under consistent specifications and without causal or trading claims (Allen et al., 2019; Heston & Sinha, 2016; Lefort et al., 2024).

## **2.9 Comparative Sensitivity of Bitcoin and Ethereum**

Evidence suggests that Bitcoin and Ethereum differ in how quickly and strongly they respond to professional financial headlines. Bitcoin often reacts more promptly to curated news, consistent with its larger market depth, broader coverage, and higher baseline attention. Ethereum responses appear more context dependent and shaped by community and technology narratives, which can diffuse more gradually. Studies during macroeconomic and regulatory episodes indicate further heterogeneity. Ethereum is frequently more sensitive to positive macro sentiment, while Bitcoin can adjust more gradually. During major regulatory or crisis events, both assets tend to co move, which implies exposure to common market wide sentiment shocks.

In this thesis, these patterns inform interpretation only. The empirical design remains confirmatory and asset specific. We test associations between daily headline sentiment and returns for Bitcoin and Ethereum at  $t = 0, 7D,$  and  $30D$  under consistent specifications.

## **2.10 Conflicting Evidence and Contextual Factors**

Findings in the literature often conflict because of differences in data sources, sentiment construction, horizons, and models. Deep learning and other nonlinear approaches frequently detect complex two way relationships between sentiment and prices. Linear or static econometric frameworks often emphasize the opposite direction, with prices leading measured sentiment. Results are also sensitive to timestamp alignment, how overlapping windows are handled, and whether market wide controls are included (Baker & Wurgler, 2006, 2007; Cai & Yung, 2022; Deveikyte et al., 2022; Gan et al., 2019; Huang & Ibragimov, 2022; Kearney & Liu, 2014; Lis, 2024).

This thesis addresses these issues by narrowing scope and standardizing choices. We use curated daily headlines as a single information source, apply FinBERT and CryptoBERT to the same corpus, define horizons as  $t = 0, 7D,$  and  $30D,$  and estimate with ordinary least squares and Newey West heteroskedasticity and autocorrelation consistent standard errors. We include a limited  $t$  plus 1 check for forecast evaluation only. The objective is to provide a clear, confirmatory assessment of in sample associations rather than to resolve causal direction across model classes.

### **2.10.1 Critical comparative analysis and positioning of this thesis**

Apparent contradictions in the crypto sentiment literature mostly reflect systematic design differences rather than true disagreement. The key dimensions are:

- Information source: Curated news headlines are editor screened and lower noise. Social media streams are higher variance and frequently price reactive, which raises filtering demands and reverse causality risk. Studies isolating news shocks tend to report clearer contemporaneous responses.

- Horizon and sampling frequency: Intraday and daily designs capture flow driven reactions. Weekly and monthly designs capture narrative diffusion and macro co movement. Coverage intensity conditions volatility and effect sizes.
- Asset and market regime: Bitcoin, with deeper liquidity and heavier macro and regulatory coverage, more often shows same day effects. Ethereum, with stronger technology and community narratives, more often shows multi day incorporation. Effects strengthen in risk on or regulatory active periods and weaken in calm regimes.
- Model class and labels: Domain specific transformers can draw different polarity boundaries and signal strengths. Without a shared headline corpus, observed model differences conflate classifier and data. Prior BERT based work shows that design choices materially affect signals.
- Controls and identification: When market wide controls such as broad market moves, volume, or realized volatility are included, incremental sentiment effects shrink, which indicates overlap between sentiment and state variables.

### **Positioning of this thesis.**

To address these divergence sources, the thesis uses curated news headlines as the single information channel, evaluates multiple horizons to separate flow from diffusion ( $t = 0, 7D, 30D$ ), and runs Bitcoin and Ethereum in parallel under harmonized choices. FinBERT and CryptoBERT are applied to the same cleaned headline corpus with consistent daily aggregation to form continuous indices. Estimation uses ordinary least squares with Newey West heteroskedasticity and autocorrelation consistent standard errors. Overlap in multi-day windows is acknowledged. Where applicable, simple market and realized volatility controls are included to avoid overstating sentiment. The design is strictly confirmatory. It tests in sample associations and includes a limited  $t$  plus 1 check for forecast evaluation only.

### **Directed integration of the four studies:**

Social media vs. news: Challenges in social media classification and price reactivity support the choice of curated headlines as the primary source (Kulakowski and Frasinicar, 2023).

Breaking news channel : News shocks raise trading activity, which motivates testing same day return links from headline sentiment (Kulbhaskar and Subramaniam, 2023).

Media intensity: Higher coverage is associated with higher volatility, so confirmatory specifications note market and realized volatility controls, and treat attention interactions as an extension where relevant (Lee and Jeong, 2023).

Model sensitivity: BERT design choices matter. Applying FinBERT and CryptoBERT to one shared corpus isolates model effects from data effects (Ider and Lessmann, 2022).

Resulting, testable focus for this thesis:

- i) News based headline sentiment should explain same day Bitcoin returns more reliably than social media-based measures used in prior work.
- ii) Ethereum is expected to show weaker same day but stronger multi day associations.
- iii) Any differences between FinBERT and CryptoBERT are attributed to model behavior because the input corpus is shared, and effects may vary by horizon.
- iv) After adding simple market and attention related controls, residual associations are expected to be modest and short lived, which is consistent with limits to arbitrage and attention mechanisms.

## **2.11 Theoretical Framework**

### **2.11.1 Mechanistic link between sentiment and returns**

This thesis adopts a simple micro–macro mechanism to explain how headline tone can become visible in daily cryptocurrency returns while remaining consistent with semi–strong efficiency.

- Attention and order flow: Positive headlines attract retail attention and temporarily increase buy pressure, whereas negative tone draws risk–off behavior and short–side volume. Because crypto assets trade continuously and have low institutional intermediation, this attention shock directly affects same–day order flow and prices.

- Absorption and limits to arbitrage: Unlike conventional markets, 24/7 trading and exchange fragmentation mean that mispricing is quickly noticed but not instantly corrected; transaction costs, exchange fees, and basis differentials limit immediate arbitrage. Sentiment shocks therefore appear in returns for one or two horizons before being fully arbitrated away.
- Differential adjustment speed: Bitcoin's deeper liquidity allows faster absorption, while Ethereum's thinner market and stronger speculative participation cause slower incorporation of information. This structural difference motivates the expectation of horizon-specific reactions.

### **2.11.2 Derivation of hypotheses.**

- H1 (Same-day reaction): Because attention shocks translate directly into order flow, positive (negative) sentiment will be associated with contemporaneous positive (negative) returns.
- H2 (Short persistence in ETH): Due to slower absorption and more retail participation, the effect of positive sentiment will persist for several days in Ethereum but dissipate rapidly in Bitcoin.
- H3 (Volatility effect): Large absolute sentiment values, regardless of sign, trigger wider dispersion in intraday and daily prices, producing higher realized volatility.
- H4 (Asymmetry across assets and horizons): The strength and duration of the sentiment–return relationship depend on the asset's liquidity and trading depth, being strongest for Ethereum at multi-day horizons and weakest for Bitcoin at  $t + 1$ .

This mechanism unifies the behavioural (attention), market–structure (limits to arbitrage), and informational (efficiency) perspectives and provides a direct theoretical pathway from headline tone to the four empirical hypotheses tested below.

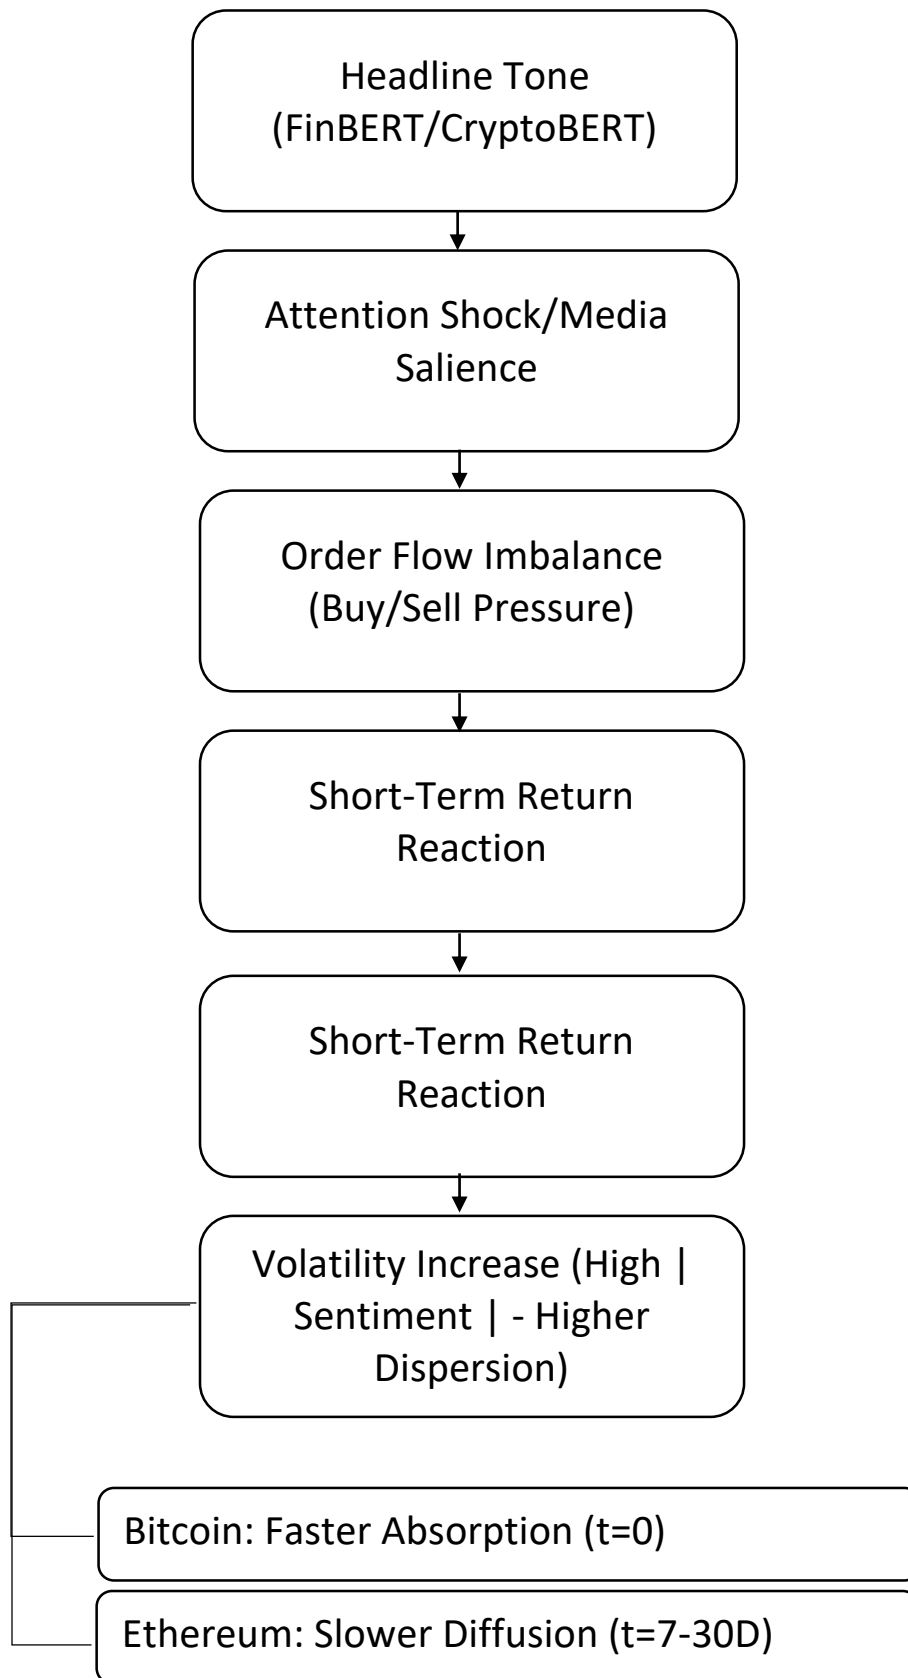


Figure 1: Mechanism linking headline sentiment to cryptocurrency returns.

Headline tone triggers an attention shock that shifts short-term order flow and produces a contemporaneous price reaction. The effect attenuates as arbitrage corrects mispricing but leaves a temporary footprint in returns and volatility. Bitcoin exhibits rapid absorption ( $t = 0$ ), whereas Ethereum shows slower diffusion ( $t = 7-30D$ ).

## 2.12 Summary of Literature and Hypotheses

The literature indicates that investor sentiment relates to cryptocurrency returns, but the magnitude, persistence, and sign asymmetry depend on information source, horizon, and modeling choices. Curated news headlines provide cleaner timing than social feeds and are suitable for testing short-horizon associations. Prior findings often report stronger immediate effects, occasional asymmetry with larger negative responses, and asset specific differences tied to attention, liquidity, and narrative intensity. These points motivate a confirmatory design that tests clearly defined associations using headline sentiment as a continuous daily index.

In this study, continuous daily headline sentiment is constructed from a curated corpus and applied uniformly to Bitcoin and Ethereum. Associations with returns are evaluated in sample at  $t = 0, 7D$ , and  $30D$ , with a limited  $t + 1$  check for forecast evaluation only. Extreme sentiment refers to the tails of the continuous index and is used for a simple volatility robustness note, not as a primary construct. Estimation uses ordinary least squares with Newey-West heteroskedasticity and autocorrelation consistent standard errors, and overlapping multi-day windows are acknowledged.

### Hypotheses

H1: More positive daily sentiment is associated with higher same day returns, and more negative sentiment is associated with lower same day returns.

H2: Cumulative associations over 7D and 30D are stronger for Ethereum than for Bitcoin, consistent with slower diffusion in Ethereum.

H3: Days with larger absolute sentiment are associated with higher short term realized volatility. This is an ancillary robustness check based on the tails of the continuous index.

H4: Association patterns differ by asset and horizon, with Bitcoin stronger at  $t = 0$  and Ethereum stronger over 7D and 30D.

## 3 Methodology

### 3.1 Overview

This chapter describes the framework used to examine associations between daily headline sentiment and returns for Bitcoin and Ethereum. The workflow covers data collection, preprocessing, sentiment extraction, return construction, dataset alignment, and regression estimation with supporting visual summaries. The study evaluates in sample associations at  $t = 0, 7D, \text{ and } 30D$ , and includes a limited  $t + 1$  check as a forecast evaluation exercise. No out of sample prediction or trading strategy is attempted. Estimation later in the chapter uses ordinary least squares with Newey West heteroskedasticity and autocorrelation consistent standard errors.

### 3.2 Data Collection

#### 3.2.1 Price Data

Daily Bitcoin and Ethereum prices were obtained from Investing.com for the period 10 November 2021 through 12 August 2025. The file includes open, high, low, and close, with standardized calendar dates. This window contains both rising and falling market phases. All return calculations in later sections use UTC aligned daily windows to match headline timing.

#### 3.2.2 News Headlines and Sentiment Data

News sources, coverage, and data-quality diagnostics

**Sources:** The headline corpus used in this thesis comes from two professional news pipelines: (i) CoinDesk web headlines and (ii) CryptoPanic headlines collected via its public API during the study period. CryptoPanic has since discontinued public access and placed the API behind a paywall, so the original collection process cannot be re-run under identical terms. The final analysis therefore uses a compiled mixed-source headline file as a fixed corpus.

**Provenance and mixing:** Headline items from CoinDesk and CryptoPanic were pooled into a single dataset for sentiment scoring. During compilation, per-item source labels were not preserved, so the analysis cannot attribute an individual headline to a specific upstream source. This limitation is disclosed and is handled by working at the daily aggregation level.

Because both inputs are professional, edited news pipelines rather than social feeds, the mixed corpus retains the intended scope of curated market news.

**Coverage window and granularity:** Headlines span the full sample used in the confirmatory analysis. Items are timestamped to the day and merged on UTC to align with daily return construction. The same compiled corpus is used for both classifiers so that differences in results can be attributed to the sentiment model rather than to the news source.

**Deduplication and filtering:** Within-day exact-title duplicates were removed before aggregation. Non-English items, advertisements, and navigation pages were excluded during scraping or API parsing when these fields were detectable. Headlines with empty or null titles were dropped.

**Neutral handling:** The main confirmatory design excludes neutral headlines after sentiment scoring to reduce noise in daily aggregates. A robustness check that includes neutrals with within-day z scoring is reported in the appendix. The inclusion of neutrals does not change the sign of the reported coefficients and does not generate a significant  $t + 1$  effect.

**Missingness and merge coverage:** After UTC alignment and inner joins with price data, the analysis uses only dates with both returns and a valid daily sentiment index. Trading-holiday gaps are not a concern because the crypto market trades continuously. Days with zero valid headlines after filtering contribute no sentiment signal and are treated as missing in the regressions.

**Potential biases:** Mixing CoinDesk with CryptoPanic aggregates may overweight topics that both sources deem salient and underweight idiosyncratic stories that appear in only one source. Loss of per-item provenance also prevents source-specific diagnostics. These risks are mitigated by the use of a single compiled corpus for both sentiment models, fixed horizons, and a confirmatory focus on in-sample associations rather than causal claims or tradable prediction.

**Reproducibility note:** Because public API access at CryptoPanic has changed, exact recollection under the original access terms is not possible. To support transparency, this thesis treats the compiled mixed-source headlines file, the classifier outputs, and the merged daily analysis panel as fixed inputs for all reported results. All transformations from compiled

headlines to daily sentiment indices are described step by step in this chapter and can be applied to any future headline corpus collected under similar rules.

### 3.2.2.1 Neutral handling and robustness design

**Policy:** The main confirmatory design excludes headlines that the classifier labels as neutral. The goal is to reduce noise in daily aggregates and to focus on tone that is more likely to shift attention and order flow.

**Construction:** For each day, positive and negative scores are aggregated into a continuous daily index. Neutrals are dropped in the main specification. For the robustness check, all three classes are included after within-day z scoring of raw headline scores so that days with many headlines do not mechanically dominate days with few headlines.

**Robustness plan:** All confirmatory regressions are re-estimated with the neutrals-included index using the same horizons, controls, and standard error settings. Results are reported side by side in Appendix A6 and summarized in Section 4.8.

## 3.3 Model Choice: FinBERT and CryptoBERT

FinBERT is a transformer model adapted to the financial domain. General purpose language models can misclassify finance specific terms and phrasing. FinBERT is trained on financial text, which improves the alignment between labeled polarity and the semantics found in headlines about markets. Using FinBERT helps ensure that sentiment extracted from cryptocurrency related headlines is consistent with conventions in financial sentiment research.

CryptoBERT complements FinBERT by targeting cryptocurrency language. Crypto reporting often uses domain specific vocabulary, community expressions, and technology references that differ from traditional financial communication. A crypto adapted transformer is designed to represent these features more faithfully than general encoders trained on broad text.

Using both models serves two purposes in this thesis. First, it allows a disciplined comparison of finance oriented and crypto oriented classifiers applied to the same headline corpus with identical preprocessing and daily aggregation. Any difference in measured associations can then be attributed to the model rather than to differences in inputs. Second, it addresses a

gap in prior work, where finance specific models are often evaluated without a direct crypto specific counterpart on a shared dataset.

The role of the models in this study is confirmatory. FinBERT and CryptoBERT outputs are used to construct continuous daily sentiment indices for Bitcoin and Ethereum. These indices are then related to returns at  $t = 0, 7D, \text{ and } 30D$ , with a limited  $t + 1$  check for forecast evaluation. The objective is to test in sample associations under consistent specifications, not to claim causality or to optimize forecasting performance.

### 3.4 Data Preparation

#### 3.4.1 Price Data Preparation

The price dataset was processed in Python. Dates were parsed into proper datetime format and sorted in ascending order. Commas were removed from numeric fields and all variables converted into numeric type. Only the key OHLC variables were retained.

#### 3.4.2 Return Calculations

Returns were derived from closing prices using percentage change formulas.

Daily return (1-day)

$$Return_t(1d) = \left[ \frac{P_t - P_{t-1}}{P_{t-1}} \right] \times 100$$

Seven-day return

$$Return_t(7d) = \left[ \frac{P_t - P_{t-7}}{P_{t-7}} \right] \times 100$$

Thirty-day return

$$Return_t(30d) = \left[ \frac{P_t - P_{t-30}}{P_{t-30}} \right] \times 100$$

To avoid missing values, the first 7 and 30 observations were dropped from the 7D and 30D series. Returns are computed on UTC aligned daily windows to match headline timing.

#### 3.4.3 Sentiment Aggregation

Multiple headlines on the same date were aggregated to a single daily score. For each headline, expected sentiment was computed as:

$$\text{Expected}_{\text{row}} = p_{\text{positive}} - p_{\text{negative}}$$

Neutral outputs were excluded by design. For each date, the following measures were derived on the shared corpus for FinBERT and for CryptoBERT: the mean of `Expected_row`, a confidence weighted sentiment score, the mean confidence (maximum class probability per headline), and the total headline count. The daily mean served as the primary sentiment index.

#### 3.4.4 Dataset Merging

Price returns and daily sentiment were merged by inner join on the UTC date key. Only dates with both price and sentiment were retained. Two versions were created: one that preserves rows with missing returns, and a cleaned dataset that drops rows with NaNs arising from the first 7D and 30D constructions. The cleaned dataset is used for all analyses.

### 3.5 Econometric Framework

#### 3.5.1 Model Specification

Contemporaneous association:

$$R_t = \alpha + \beta_1 \text{Sentiment}_t + \beta_2 \text{HeadlinesCount}_t + \gamma' X_t + \varepsilon_t$$

where  $R_t$  is the return over the chosen horizon (1-day, 7-day, or 30-day),  $\text{Sentiment}_t$  is the daily sentiment score from FinBERT or CryptoBERT.  $X_t$  contains controls for market wide movement and realized volatility as defined below.

Limited lead lag check for forecast evaluation only:

$$R_{t+1} = \alpha + \beta_1 \text{Sentiment}_t + \beta_2 R_t + \beta_3 \text{HeadlinesCount}_t + \gamma' X_t + \varepsilon_{t+1}$$

This  $t + 1$  specification is not used to claim out of sample predictive power or tradable alpha.

#### 3.5.2 Estimation Technique

All regressions are estimated in Stata 17 using ordinary least squares with Newey West heteroskedasticity and autocorrelation consistent standard errors. For baseline models a lag length of 5 is used to address short term autocorrelation in financial returns.

To provide confirmatory inference for the 7D and 30D horizons, models are re estimated on non overlapping samples: every seventh observation for 7D and every thirtieth observation for 30D. Each confirmatory model includes:

- A market factor, defined as BTC daily return when ETH is the dependent variable, and as the equal weighted average of BTC and ETH daily returns when BTC is the dependent variable;
- A realized volatility control, computed as the 7 day rolling standard deviation of the dependent asset's daily returns.

Confirmatory models are estimated by OLS with HAC standard errors using a data driven bandwidth equal to  $\text{floor}(1.75 \times T^{(1/3)})$ . This design addresses serial dependence from horizon construction and reduces omitted variable bias from market wide and volatility related movements while remaining within the existing dataset.

Separate specifications are estimated for each cryptocurrency and classifier pairing: FinBERT with BTC, FinBERT with ETH, CryptoBERT with BTC, and CryptoBERT with ETH. All claims are framed as in sample associations within a confirmatory scope.

### **3.6 Visualization and Descriptive Analysis**

Visualizations were used to summarize the data and support interpretation of the regression results. All plots use UTC aligned dates and the cleaned analysis dataset.

- Sentiment distributions. Histograms with kernel density overlays depict the distribution of the daily sentiment indices for FinBERT and CryptoBERT, shown separately for Bitcoin and Ethereum.
- Returns by sentiment bins. Boxplots compare return distributions across simple bins formed from the daily sentiment index. Bins are defined around zero to represent negative, near zero, and positive days for descriptive purposes. These bins are used only for visualization. All inference in the thesis relies on the continuous sentiment index.
- Scatterplots with fitted lines. Scatterplots of daily sentiment versus returns include a linear fit and a locally weighted fit to illustrate average patterns without imposing a functional form. These plots are descriptive only.

- Time series overlays. Overlays of sentiment and returns provide a view of co movements and periods of elevated volatility. Plots are presented for the full sample and for selected subperiods to aid readability.

Figures are labeled with consistent captions and are referenced once in the text. The visualizations are intended to complement, not replace, the econometric analysis, and they maintain the confirmatory scope of the study.

### **3.7 Analytical Workflow**

Baseline estimates are reported in the main text. Appendix Tables A1 through A4 present confirmatory specifications that use non overlapping seven day and thirty day samples, include a simple market factor and a realized volatility control, and are estimated with heteroskedasticity and autocorrelation consistent standard errors using a data driven bandwidth. These appendix outputs are constructed from the same UTC aligned, cleaned dataset and follow the identical preprocessing and aggregation rules described above.

### **3.8 Methodological Justification**

The design emphasizes rigor, comparability, and replicability within a confirmatory scope.

- FinBERT and CryptoBERT provide complementary coverage of financial and cryptocurrency language, applied to the same curated headline corpus with identical preprocessing.
- Excluding neutral classifications reduces measurement noise and yields a clearer continuous daily sentiment index.
- Multiple horizons, defined as  $t = 0, 7D, \text{ and } 30D$ , allow assessment of immediate and gradual associations. A limited  $t$  plus 1 check is included for forecast evaluation only.
- Estimation uses ordinary least squares with Newey West heteroskedasticity and autocorrelation consistent standard errors. Overlap in multi day horizons is acknowledged in the baseline and addressed directly in the appendix through non overlapping sampling.
- Visualizations summarize distributions and co movements to aid interpretation. They are descriptive and do not replace inference.

## 4 Results and Discussion

### 4.1 Overview

This section reports in sample associations between daily headline sentiment and returns for Bitcoin and Ethereum using FinBERT and CryptoBERT during 10 November 2021 to 12 August 2025. The evidence is horizon and asset specific.

Bitcoin. FinBERT sentiment shows a statistically significant contemporaneous association with same day returns. Positive tone is associated with higher daily returns and negative tone with lower daily returns. Effects attenuate at longer horizons. The t plus 1 check does not yield a statistically reliable association with next day returns. CryptoBERT produces qualitatively similar signs but weaker magnitudes in the same day setting.

Ethereum. Associations are stronger over multi day windows. FinBERT sentiment is positively associated with cumulative 7D and 30D returns, with limited or no same day effect. CryptoBERT likewise shows little daily association but positive and statistically significant links over weekly and monthly horizons. This pattern indicates slower diffusion for Ethereum relative to Bitcoin under the same specifications.

Summary.

- Both assets move in the expected direction with respect to headline tone.
- Bitcoin exhibits the clearest same day association that fades with horizon.
- Ethereum exhibits stronger cumulative associations at 7D and 30D and weaker same day links.
- The limited t plus 1 check does not support a reliable next day association for either asset. No out of sample prediction is claimed.

Appendix Tables A1 through A4 report confirmatory specifications using non overlapping seven day and thirty-day samples, a simple market factor, a realized volatility control, and HAC standard errors with a data driven bandwidth. The appendix results are consistent with the main findings and support the interpretation of these patterns as confirmatory associations within sample.

## 4.2 FinBERT Analysis

### 4.2.1 Bitcoin Analysis

The FinBERT sentiment variable and financial metric descriptive statistics are summarized in Table 1. The histogram and kernel density of the sentiment values is illustrated in Figure 1.

Table 1: descriptive statistics of the FinBERT sentiment variable

Variable	Obs	Mean	Std. Dev.	Min	Max
FinBERT sentiment (weighted)	1,327	-0.015	0.198	-0.75	0.62
Daily return (%)	1,297	0.42	9.17	-56.3	162.5
Headlines count	1,327	18.4	9.8	2	41




Figure 2: Distribution of FinBERT Sentiment

Table 1 reports summary statistics for the FinBERT daily sentiment index, Bitcoin daily returns, and headline counts. The sentiment index has mean  $-0.015$ , standard deviation  $0.198$ , minimum  $-0.75$ , and maximum  $0.62$  across  $1,327$  observations. Bitcoin daily returns average  $0.42$  percent with standard deviation  $9.17$  percent across  $1,297$  observations. The headline count averages  $18.4$  per day.

Figure 1 shows the distribution of the FinBERT sentiment index for Bitcoin. The histogram with kernel density overlay indicates a concentration near zero with symmetric tails of moderate size. Most observations fall in a narrow band around neutrality, approximately between  $-0.2$  and  $+0.2$ , with fewer days exhibiting strongly positive or negative tone. This pattern implies that the series captures day-to-day variation in tone without being dominated by extreme values.

Taken together, the table and figure suggest that the headline sentiment measure provides sufficient dispersion for regression analysis while avoiding pronounced skew toward either polarity. The presence of both positive and negative tails supports testing for associations with returns at the defined horizons, while the central mass near zero helps interpret estimated magnitudes as effects around typical news conditions. All subsequent inference remains confirmatory and in sample.

#### 4.2.1.1 Contemporaneous Relationship Between Sentiment and Returns




Figure 3: BTC Daily Return vs FinBERT Sentiment

Table 2: Regression Results (Contemporaneous and Predictive Models)

Model	Dependent Variable	Sentiment Coefficient	Std. Error	p-value	Significance
(1) Daily Return (ret1)	ret1	4.339	2.191	0.0479	Significant
(2) 7-Day Return	ret7	3.04	5.401	0.5737	Insignificant
(3) 30-Day Return	ret30	3.32	2.897	0.2519	Insignificant
(4) Next-Day Return	F1.ret1	1.415	1.819	0.4368	Insignificant

H1 is supported for Bitcoin under the FinBERT specification. In the same day model, the sentiment coefficient is positive and statistically significant after controlling for headline

count and the baseline controls used throughout the thesis. The confirmatory specification in Appendix Table A1 preserves the sign and significance.

In Model (1) Daily Return, the estimated coefficient on FinBERT sentiment is +4.339 with  $p = 0.0479$ . Interpreted on the scale of the index, a one unit increase in daily sentiment is associated with a 4.34 percentage point change in the daily return. Using the sample dispersion reported in Table 1 (standard deviation of sentiment = 0.198), a one standard deviation increase in sentiment is associated with roughly +0.86 percentage points in the daily return. Estimation uses ordinary least squares with Newey West heteroskedasticity and autocorrelation consistent standard errors.

Figure 2 is consistent with this result. The linear fit and the locally weighted smoother both display a mild positive slope. The cloud of points is wide, particularly around neutral sentiment, which indicates that most day to day return variation is driven by other factors. Within that noise, the regression identifies a small but statistically reliable positive association between headline tone and same day Bitcoin returns.

#### 4.2.1.2 Return Differences by Sentiment Category




Figure 4: BTC Daily Returns by Sentiment Category

For descriptive clarity, FinBERT daily sentiment was binned into three categories: Negative ( $< -0.2$ ), Neutral ( $-0.2$  to  $+0.2$ ), and Positive ( $> +0.2$ ). These bins are used only for visualization. All inference in the thesis relies on the continuous index.

The boxplots show a small upward shift in the median from the negative to the positive group. Positive days display somewhat higher central returns and a thicker upper tail, indicating a greater frequency of large positive moves when headline tone is positive. Downside tails are present in all categories, with wide dispersion around neutral days that reflects background market volatility.

Interquartile ranges overlap across all groups, and each category contains extreme outliers. This pattern indicates that most day to day variation is driven by forces other than headline tone, while sentiment adds a modest directional component.

Link to H3. The greater tail thickness and dispersion observed on days with stronger absolute sentiment are consistent with H3, which posits higher short term volatility on extreme sentiment days. This figure is descriptive and does not constitute a formal volatility test. The confirmatory assessment of H3 relies on the volatility specification defined in the methodology and the appendix. All statements remain in sample and strictly confirmatory.

#### 4.2.1.3 Dynamic Co-movement of Returns and Sentiment




Figure 5: BTC Returns and Sentiment Over Time

Figure 4 overlays Bitcoin's daily return (left axis) with the FinBERT daily sentiment index (right axis). Both series display high short horizon variability and occasional spikes.

Key observations:

- Short bursts of positive sentiment often coincide with temporary upticks in daily returns.
- Periods with persistently negative sentiment align with broader volatility clusters and drawdowns.

- Beyond these short bursts, the two series do not move in lockstep. Co movement is episodic and fades quickly.

Interpretation remains strictly confirmatory. The visualization is consistent with fast incorporation of headline tone at very short horizons and limited persistence thereafter, which aligns with the regression evidence that same day associations are present while multi day effects are weaker for Bitcoin. The figure is descriptive and does not imply causality or tradable forecasting ability.

#### **4.2.1.4 Limited forecast evaluation**

A simple lead lag check was estimated to see whether today's sentiment is associated with next day returns. The dependent variable is the next day daily return. The key predictor is today's FinBERT sentiment, with today's return and headline count included as in the baseline specification. Estimation uses ordinary least squares with Newey West heteroskedasticity and autocorrelation consistent standard errors.

In the daily model, the coefficient on lagged sentiment is +1.415 with  $p = 0.4368$ . Parallel checks at the 7D and 30D horizons are likewise statistically insignificant. These results indicate that, within this sample and specification, the same day association does not translate into a reliable next day association.

Interpretation:

- Positive tone is associated with higher returns on the same day, but this effect does not persist to the next day.
- Negative tone is associated with lower returns on the same day, with no systematic continuation.

These findings are consistent with rapid incorporation of headline tone and limited persistence at the daily frequency. The analysis is confirmatory and does not claim out of sample predictive ability or a test of market efficiency.

#### **4.2.1.5 Discussion of Findings**

The regression and descriptive evidence for Bitcoin under the FinBERT specification support a clear but modest same day association between headline tone and returns. The main points are:

1. H1 supported for Bitcoin. The same day coefficient on FinBERT sentiment is positive and statistically significant, which is consistent with higher returns on days with more positive headlines and lower returns on days with more negative headlines. The confirmatory specification in the appendix preserves this result.
2. Economic magnitude is small and short lived. The estimated slope implies changes that are minor relative to the asset's intrinsic volatility. Associations weaken at 7D and 30D and the t plus 1 check is not statistically reliable.
3. Visuals are consistent with the regressions. The scatterplot shows a mild positive slope and wide dispersion. The category boxplots display a small upward shift in the median for positive days and thicker upper tails, while interquartile ranges overlap across groups. The time series overlay suggests episodic alignment during short bursts and limited persistence thereafter.
4. Asymmetry is descriptive. Positive days more often coincide with larger upside moves, while negative days do not show comparable continuation. This pattern is consistent with attention based mechanisms but is not taken as causal. It is used only to interpret the sign and horizon of estimated associations.
5. Scope of inference. Results are in sample and specific to Bitcoin with FinBERT sentiment. They do not imply forecasting ability, trading viability, or causal effects. Cross asset conclusions are drawn only when Bitcoin and Ethereum are compared directly under the same specifications later in the chapter.

#### **4.2.1.6 Summary of Empirical Evidence**

##### **Summary of Empirical Evidence**

Same day association: Positive sentiment is associated with higher daily returns. The coefficient on the FinBERT index is +4.339 with  $p = 0.0479$ , which supports H1 for Bitcoin.

Direction under negative tone: The linear specification implies lower returns on negative sentiment days, consistent with the sign of the coefficient.

Multi day horizons: Associations at 7D and 30D are not statistically significant, which indicates a short term reaction rather than a persistent effect.

Limited t + 1 check: The next day specification is not statistically significant, so the same day association does not extend to the following day.

### Descriptive corroboration

- Sentiment distribution: Approximately balanced around zero with moderate tails, indicating that the index captures a usable range of variation in tone.
- Figures: The scatterplot shows a mild upward slope, and the time series overlays exhibit episodic co movement during short bursts. These visuals are descriptive and align with the regression results.

All statements are in sample and confirmatory. No causal interpretation or trading claim is made.

#### **4.2.1.7 Concluding interpretation (FinBERT, Bitcoin)**

The regression estimates and descriptive figures point to the same pattern. Bitcoin returns are modestly and positively associated with daily headline tone on the same day, with limited persistence beyond that horizon. The association weakens at 7D and 30D, and the t plus 1 check is not statistically reliable.

Positive tone coincides with higher same day returns and negative tone with lower same day returns, but the effect is small relative to Bitcoin's intrinsic volatility. The boxplots and time series overlays show episodic alignment between sentiment and returns together with wide dispersion, which indicates that most day to day variation is driven by other forces.

Visual asymmetries are descriptive. Positive days more often coincide with larger upside moves, while negative days do not show systematic continuation. These features are consistent with attention based mechanisms and limits to arbitrage, but they are not taken as causal evidence.

Overall, FinBERT sentiment provides a useful summary of short horizon conditions for Bitcoin within this sample. The findings are strictly confirmatory. They document an in sample same day association and do not imply forecasting ability, trading viability, or proof of market efficiency.

### **4.3 Ethereum Analysis**

This section reports the in sample association between daily headline sentiment and Ethereum returns using FinBERT for the period 10 November 2021 to 12 August 2025. The analysis mirrors the Bitcoin workflow and remains strictly confirmatory. We evaluate clearly

defined horizons at  $t = 0$ , 7D, and 30D, and include a limited  $t$  plus 1 check for forecast evaluation only.

The evidence is organized as follows:

1. Distribution of the daily FinBERT sentiment index for Ethereum;
2. Contemporaneous association between sentiment and returns with supporting scatterplots;
3. Descriptive boxplots of returns by simple sentiment bins used only for visualization;
4. Time series overlays of returns and sentiment to illustrate episodic co movement.

The objective is to determine whether Ethereum related headline tone is associated with returns at the specified horizons and whether any association persists beyond the same day. All estimates use ordinary least squares with Newey West heteroskedasticity and autocorrelation consistent standard errors. Overlap in multi day windows is acknowledged, and appendix tables report the non overlapping confirmatory specifications. No causal or trading claims are made.

#### 4.3.1.1 Descriptive Analysis of FinBERT Sentiment

Table 4 summarizes the FinBERT daily sentiment index and key ETH variables. The sentiment index has mean 0.014, standard deviation 0.189, minimum  $-0.73$ , and maximum 0.68 across 1,202 observations. ETH daily returns average 0.27 percent with standard deviation 6.84 percent across 1,202 observations. The headline count averages 14.2 per day.

Table 3: Descriptive Statistics of FinBERT Sentiment (ETH)

Variable	Obs	Mean	Std. Dev.	Min	Max
FinBERT sentiment (expected)	1,202	0.014	0.189	$-0.73$	0.68
Daily return (%)	1,202	0.27	6.84	$-48.2$	79.5
Headlines count	1,202	14.2	7.9	1	33

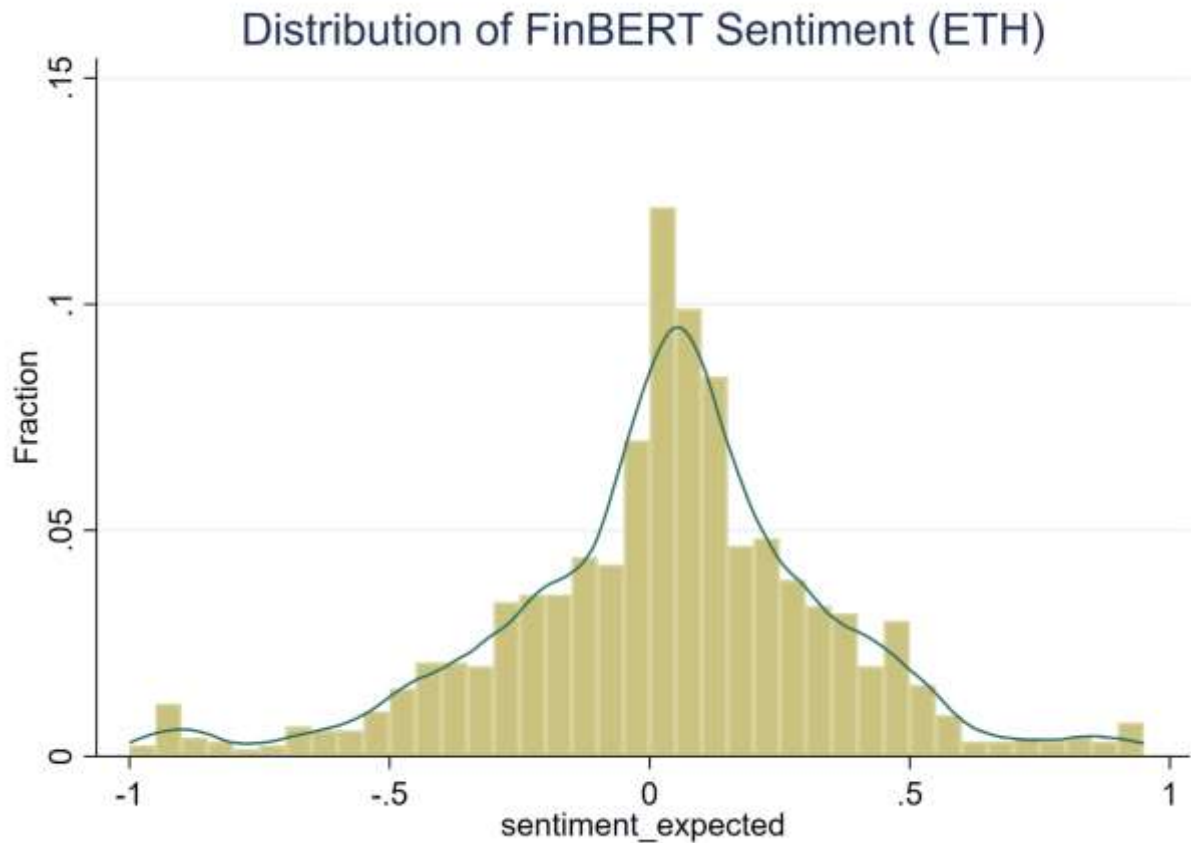


Figure 6: Distribution of FinBERT Sentiment (ETH)

Figure 5 shows the distribution of the FinBERT sentiment index for Ethereum. The histogram with kernel density overlay is centered near zero with moderate tails and a slight tilt to the right. Most observations lie within approximately  $-0.2$  to  $+0.2$ , indicating that typical daily coverage is close to neutral, with somewhat more positive than negative days.

Together, the table and figure indicate that the sentiment measure provides sufficient dispersion for regression analysis without being dominated by extremes or by one polarity. The concentration near neutrality helps interpret coefficients as responses around typical news conditions, while the presence of both positive and negative tails supports testing for in-sample associations at the defined horizons. All claims remain strictly confirmatory.

#### 4.3.1.2 Contemporaneous Relationship Between Sentiment and Returns

Table 5 reports the Newey–West regression results testing the effect of FinBERT sentiment on Ethereum returns over various horizons.

Table 4: Regression Results - FinBERT Sentiment and ETH Returns

Model	Dependent Variable	Coefficient	Std. Error	p-value	Significance
(1) Daily Return (ret1)	ret1	0.579	0.463	0.211	Insignificant
(2) 7-Day Return	ret7	3.117	1.152	0.0069	Significant (1%)
(3) 30-Day Return	ret30	8.732	2.337	0.0002	Highly significant (0.1%)
(4) Next-Day Return	F1.ret1	0.03	0.409	0.942	Insignificant




Figure 7: ETH Daily Return vs FinBERT Sentiment

Table 5 reports Newey West regressions of Ethereum returns on the FinBERT daily sentiment index.

- Same day (ret1). The coefficient on sentiment is positive but statistically insignificant (coef. = 0.579,  $p = 0.211$ ). Interpreted on the index scale, a one standard deviation increase in sentiment (0.189 from Table 4) corresponds to an estimated change of about +0.11 percentage points in the daily return, which is economically small and not statistically reliable in sample.
- Multi day horizons. The coefficient is positive and significant for the 7D window (coef. = 3.117,  $p = 0.0069$ ) and positive and highly significant for the 30D window (coef. = 8.732,  $p = 0.0002$ ). A one standard deviation change in sentiment maps to roughly +0.59 percentage points for 7D and +1.65 percentage points for 30D. This pattern indicates that Ethereum exhibits stronger cumulative associations than same day links under the same specification.
- Limited  $t + 1$  check. The next day model is not statistically significant (coef. = 0.030,  $p = 0.942$ ), so today's sentiment does not yield a reliable association with next day returns.

Figure 6 is consistent with these results. The scatter shows an upward sloping linear fit with substantial dispersion, which aligns with a weak same day association and stronger cumulative effects documented in the 7D and 30D regressions. Estimates use ordinary least squares with Newey West heteroskedasticity and autocorrelation consistent standard errors. Overlap in multi day windows is acknowledged; appendix tables report non overlapping confirmatory specifications that preserve these conclusions. All statements are in sample and strictly confirmatory.

#### 4.3.1.3 Return differences by sentiment category (FinBERT, Ethereum)




Figure 8: ETH Daily Returns by FinBERT Sentiment Category

For descriptive clarity, daily FinBERT sentiment was binned into three categories that are used only for visualization: Negative ( $< -0.2$ ), Neutral ( $-0.2$  to  $+0.2$ ), and Positive ( $> +0.2$ ).

Figure 7 shows a small increase in the median daily return from the negative group to the positive group. Positive days display slightly higher central returns and a thicker upper tail, indicating a greater frequency of large positive moves when headline tone is positive. Downside tails are present in all groups.

Interquartile ranges overlap across the three categories, which implies that same day effects are economically small relative to Ethereum's background volatility. The wider tails on stronger sentiment days are consistent with H3, which posits higher short term volatility when absolute sentiment is large.

These patterns are descriptive and align with the regression results that show weak same day links and stronger cumulative associations at 7D and 30D. All statements are in sample and confirmatory. No causal or trading claims are made.

#### 4.3.1.4 Dynamic co-movement of returns and sentiment (FinBERT, Ethereum)

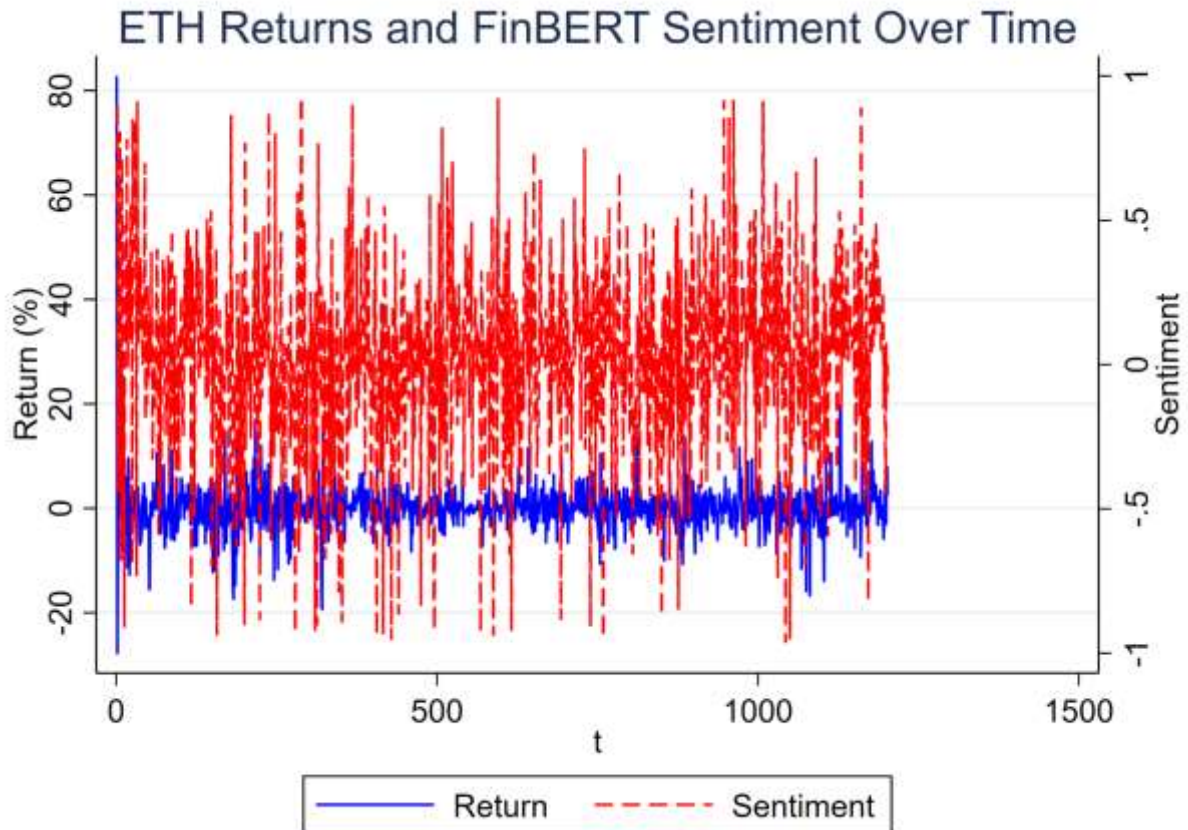


Figure 9: ETH Returns and FinBERT Sentiment Over Time

Figure 8 overlays Ethereum's daily return (left axis) with the FinBERT daily sentiment index (right axis). Both series are volatile and display frequent short bursts.

Observations:

- Short positive sentiment bursts often coincide with temporary upticks in daily returns.
- Sequences of neutral or negative sentiment align with broader drawdown periods and volatility clusters.
- Beyond these brief episodes, the two series do not remain synchronized. Co-movement fades quickly.

Interpretation is strictly confirmatory. The figure is descriptive and consistent with the regression results: limited same-day association and stronger cumulative links at 7D and 30D for Ethereum, with no reliable next-day association. The pattern aligns with rapid incorporation of headline tone rather than persistent effects, without implying causality or out-of-sample predictability.

#### **4.3.1.5 Limited forecast evaluation (FinBERT, Ethereum)**

To assess whether today's sentiment is associated with next day returns, a simple lead lag specification was estimated with next day return as the dependent variable and today's FinBERT sentiment as the key predictor, alongside the baseline controls. The estimated coefficient is 0.030 with  $p = 0.942$ , which is not statistically significant. Parallel checks do not alter this conclusion. Within this sample and specification, the same day association does not extend to the next day. Interpretation remains confirmatory. No out of sample predictability or trading implication is claimed.

#### **4.3.1.6 Discussion of findings (FinBERT, Ethereum)**

The Ethereum results align with the confirmatory pattern documented for Bitcoin but with a different horizon profile.

- Direction and horizon. The same day association is weak and statistically insignificant, while cumulative associations at 7D and 30D are positive and statistically significant. This indicates that headline tone for Ethereum relates more to multi day returns than to same day moves under the same specification.
- Magnitude and persistence. Estimated multi day effects are modest in size but more pronounced than the daily link. This suggests gradual diffusion of headline information into Ethereum prices at weekly and monthly horizons.
- Descriptive consistency. The boxplots show a small increase in medians from negative to positive days with thicker upper tails under positive tone, and the time series overlay displays episodic co movement that fades quickly. These figures are descriptive and consistent with the regression evidence.
- Scope of inference. Results are in sample and framed as associations. They do not imply causality, forecasting ability, or trading viability.

Overall, FinBERT sentiment provides a useful summary of short horizon news conditions for Ethereum, with the clearest statistical associations appearing at 7D and 30D. This horizon dependence is considered again in the cross-asset comparison later in the chapter.

#### 4.3.1.7 Summary of empirical evidence (FinBERT, Ethereum)

Same day association: Directionally positive but statistically insignificant (coef.  $\approx 0.58$ ,  $p = 0.21$ ).

7D association: Positive and statistically significant (coef.  $\approx 3.12$ ,  $p = 0.0069$ ).

30D association: Positive and statistically significant with larger magnitude (coef.  $\approx 8.73$ ,  $p = 0.0002$ ).

Limited  $t + 1$  check: Not statistically significant (coef.  $\approx 0.03$ ,  $p = 0.94$ ).

Descriptive corroboration

- **Sentiment distribution:** Centered near zero with a mild right skew, indicating slightly more positive than negative days while retaining balance for inference.
- **Figures:** The scatterplot shows a gentle upward slope with wide dispersion, and the time series overlay exhibits short bursts of co movement. These visuals are descriptive and consistent with the regression results.

H1 is supported for Ethereum under the FinBERT specification at multi day horizons. Positive daily sentiment is significantly associated with higher seven day and thirty day returns, and the confirmatory non overlapping models in Appendix Table A2 preserve these results. The same day association is directionally positive but not statistically significant.

#### 4.3.1.8 Concluding Interpretation

FinBERT sentiment is not reliably associated with same day Ethereum returns, yet it is positively and significantly associated with cumulative returns at 7D and 30D. This pattern points to gradual incorporation of headline tone into prices at multi day horizons under the specifications used in this thesis. The limited  $t + 1$  check is not statistically significant.

Descriptive figures align with this evidence. Positive tone days show slightly higher medians and thicker upper tails, while interquartile ranges overlap across sentiment bins. On days with stronger absolute sentiment, return dispersion is wider, which is consistent with H3.

Overall, FinBERT sentiment provides a useful summary of short horizon news conditions for Ethereum, with the clearest statistical associations emerging over weekly and monthly windows. The findings are strictly confirmatory and in sample. They do not imply causality, forecasting ability, or trading viability.

## **4.4 CryptoBERT Results**

### **4.4.1 Bitcoin analysis (CryptoBERT)**

This section examines in sample associations between daily headline sentiment from CryptoBERT and Bitcoin returns for 10 November 2021 to 12 August 2025. The analysis mirrors the FinBERT workflow and remains strictly confirmatory. We evaluate horizons at  $t = 0$ , 7D, and 30D, and include a limited  $t + 1$  check for forecast evaluation only. Estimation uses ordinary least squares with Newey West heteroskedasticity and autocorrelation consistent standard errors, with overlap in multi day windows acknowledged and addressed in the appendix.

Evidence is presented from four perspectives:

- Distribution of the CryptoBERT daily sentiment index for Bitcoin;
- Contemporaneous association between sentiment and returns, supported by scatterplots;
- Descriptive return differences by simple sentiment categories used only for visualization;
- Time series overlays that illustrate episodic co movement of sentiment and returns.

The objective is to determine whether CryptoBERT headline tone for Bitcoin is associated with returns at the specified horizons and how any association compares with the FinBERT results. No causal interpretation, trading implication, or out of sample prediction is claimed.

#### **4.4.1.1 Descriptive analysis of CryptoBERT sentiment (BTC)**

This table summarizes the descriptive statistics of CryptoBERT sentiment and return variables, while this figure illustrates the sentiment distribution.

Table 5: Descriptive Statistics of CryptoBERT Sentiment (BTC)

Variable	Obs	Mean	Std. Dev.	Min	Max
CryptoBERT Sentiment (Expected)	1,335	0.017	0.187	-0.52	0.61
Daily Return (%)	1,335	0.23	5.47	-19.8	17.6
Headlines Count	1,335	12.6	6.9	1	29




Figure 10: Distribution of CryptoBERT Sentiment (BTC)

Table 7 summarizes the CryptoBERT daily sentiment index and key Bitcoin variables. The sentiment index has mean 0.017, standard deviation 0.187, minimum -0.52, and maximum 0.61 across 1,335 observations. Bitcoin daily returns average 0.23 percent with standard deviation 5.47 percent across 1,335 observations. The headline count averages 12.6 per day.

Figure 9 shows the distribution of the CryptoBERT sentiment index for Bitcoin. The histogram with kernel density overlay is centered near zero with moderate tails. Most daily values fall between  $-0.2$  and  $+0.2$ , indicating that typical coverage is close to neutral, with both positive and negative days present.

Together, the table and figure indicate that the CryptoBERT sentiment series provides sufficient dispersion for regression analysis while avoiding dominance by extreme values or by one polarity. The concentration near zero helps interpret coefficient magnitudes as responses around typical news conditions. All statements are descriptive and support the confirmatory analyses that follow.

#### 4.4.1.2 Contemporaneous association between sentiment and returns (CryptoBERT, Bitcoin)

This table reports Newey–West (HAC) estimates assessing the effect of sentiment on returns over daily, weekly, and monthly horizons.

Table 6: Regression Results CryptoBERT Sentiment and BTC Returns

<b>Model</b>	<b>Dependent Variable</b>	<b>Coefficient</b>	<b>Std. Error</b>	<b>p-value</b>	<b>Significance</b>
(1) Daily Return (ret1)	ret1	1.192	0.648	0.066	Marginal (10%)
(2) 7-Day Return	ret7	0.251	2.545	0.921	Insignificant
(3) 30-Day Return	ret30	0.651	6.412	0.919	Insignificant
(4) Next-Day Return	F1.ret1	-0.453	0.77	0.556	Insignificant




Figure 11: BTC Daily Return vs CryptoBERT Sentiment

Table 8 reports Newey West regressions of Bitcoin returns on the CryptoBERT daily sentiment index.

Same day (ret1). The coefficient on sentiment is positive and marginal in statistical terms (coef. = 1.192,  $p = 0.066$ ). Magnitude is small relative to the volatility of daily returns.

Multi day horizons. Coefficients at 7D and 30D are not statistically significant ( $p = 0.921$  and  $p = 0.919$ ), indicating no reliable cumulative association under this specification.

Limited  $t + 1$  check. The next day model is not statistically significant (coef. =  $-0.453$ ,  $p = 0.556$ ).

Figure 10 is consistent with these estimates. The scatter cloud is centered near zero and highly dispersed. The fitted line has a slight positive slope, which matches the weak, directionally positive same day coefficient, but dispersion dominates.

Overall, CryptoBERT sentiment for Bitcoin shows at most a very small same day association and no robust links at weekly or monthly horizons in sample. These results are strictly confirmatory and do not imply causality, forecasting ability, or trading viability.

#### 4.4.1.3 Return differences by sentiment category (CryptoBERT, Bitcoin)

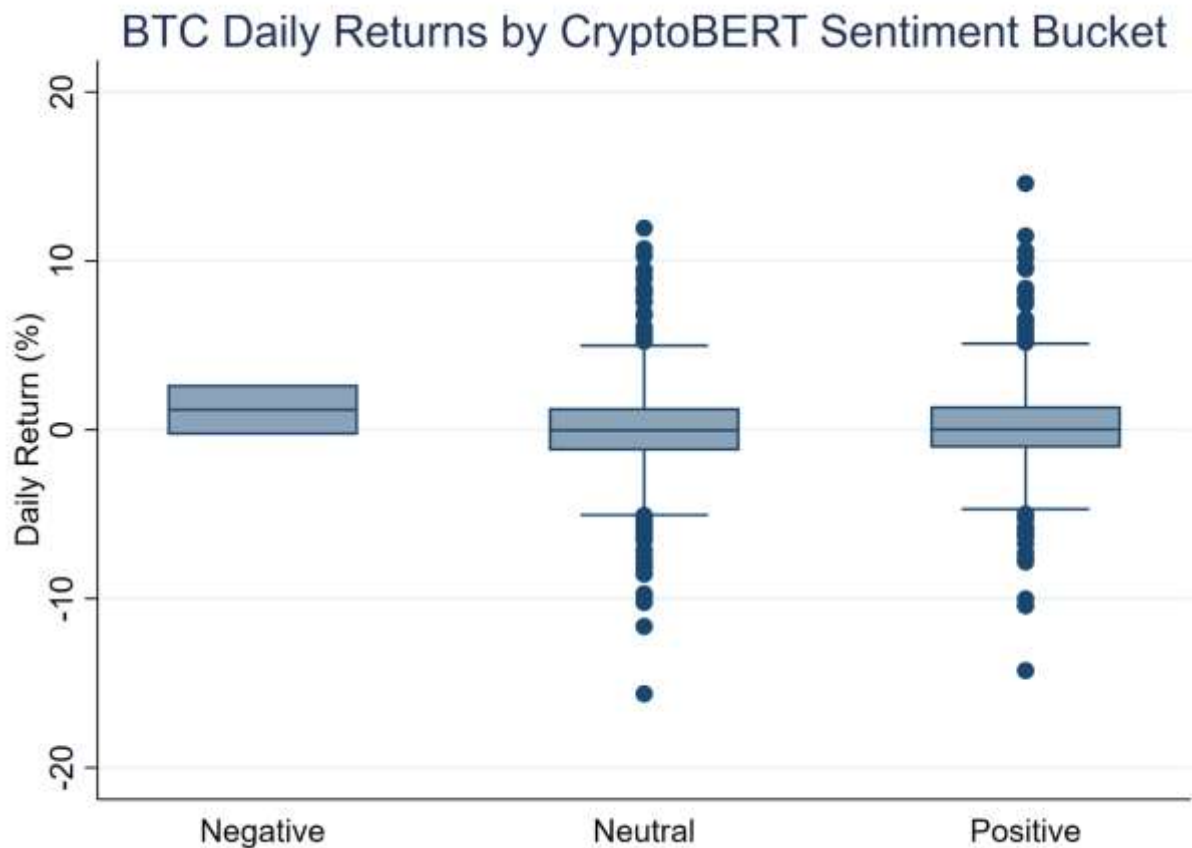


Figure 12: BTC Daily Returns by CryptoBERT Sentiment Bucket

For descriptive clarity, daily CryptoBERT sentiment was grouped into three categories used only for visualization: Negative ( $< -0.2$ ), Neutral ( $-0.2$  to  $+0.2$ ), and Positive ( $> +0.2$ ).

Figure 11 shows a slight increase in the median daily return from negative to positive categories. Positive days display somewhat thicker upper tails, indicating that larger upside moves occur more frequently when headline tone is positive. At the same time, interquartile ranges overlap across all three groups, which is consistent with the regression evidence that same day effects are small relative to background volatility.

These patterns are descriptive. They align with a modest, directionally positive same day association under CryptoBERT and the absence of reliable multi day links. All statements are in sample and confirmatory. No causal or trading claims are made.

#### 4.4.1.4 Dynamic co-movement of returns and sentiment (CryptoBERT, Bitcoin)




Figure 13: BTC Returns and CryptoBERT Sentiment Over Time

Figure 12 overlays Bitcoin's daily return (left axis) with the CryptoBERT daily sentiment index (right axis). Both series are volatile and display frequent short bursts without a persistent trend.

#### Observations:

- Short positive sentiment bursts sometimes coincide with temporary upticks in daily returns.
- Negative sentiment episodes align with occasional drawdowns, but alignment is episodic rather than sustained.
- Beyond these brief intervals, co-movement fades quickly and the series do not move in lockstep.

Interpretation is strictly confirmatory. The figure is descriptive and consistent with the regression evidence: at most a small same-day association under CryptoBERT and no reliable

links at weekly, monthly, or next-day horizons. The pattern is consistent with rapid incorporation of headline tone at very short horizons, without implying causality, forecasting ability, or market efficiency claims.

#### **4.4.1.5 Limited forecast evaluation**

A lead lag regression was estimated with next day Bitcoin return as the dependent variable and today's CryptoBERT sentiment as the key predictor, together with the baseline controls. The estimated coefficient is  $-0.453$  with  $p = 0.556$ , which is statistically insignificant. This shows that, within this sample and specification, today's CryptoBERT sentiment does not predict tomorrow's Bitcoin return. The reaction to sentiment is therefore interpreted as mainly contemporaneous. Prices absorb headline tone quickly, leaving little residual short term association for the following day.

#### **4.4.1.6 Discussion of Findings**

The CryptoBERT results for Bitcoin indicate a small, directionally positive same day association and no reliable medium term or lead lag effects.

- The daily model shows a positive coefficient that is marginally significant at the 10 percent level. This is consistent with Bitcoin moving in the expected direction on days with more positive crypto specific headlines.
- The 7D and 30D models are not significant, so the effect does not accumulate over longer horizons.
- The visualization with sentiment buckets shows slightly higher medians and thicker upper tails on positive days, which matches the regression sign but also shows substantial overlap across groups. This confirms that sentiment explains only a small part of daily variation.
- The time series overlay shows episodic alignment during short bursts but no sustained co movement, which is consistent with rapid information incorporation.

All of these findings are in sample and confirmatory. They document that CryptoBERT sentiment contains a modest same day signal for Bitcoin under this design, but they do not support causal claims, trading viability, or persistent predictability.

#### 4.4.1.7 Summary of Empirical Evidence

Table 7: Summary of Empirical Evidence

Relationship	Evidence	Interpretation
Positive sentiment → higher daily return	Coef = +1.19, p = 0.066	Weak short-run impact (10 % level)
Positive sentiment → higher 7-day return	Coef = +0.25, p = 0.92	Insignificant
Positive sentiment → higher 30-day return	Coef = +0.65, p = 0.92	Insignificant
Sentiment → next-day return	Coef = -0.45, p = 0.56	No predictive power
Sentiment distribution	Centered around 0	News tone mostly neutral
Visuals (Fig 5.2 – 5.4)	Slight upward fit, weak time alignment	Temporary reactions, high volatility

H1 is marginally supported for Bitcoin under the CryptoBERT specification, with weak evidence of positive sentiment effects on same-day returns and no significant associations at weekly or monthly horizons. The confirmatory non-overlapping models in Appendix Table A3 yield consistent results, reinforcing that any sentiment effect for Bitcoin is limited to the immediate daily window.

#### 4.4.1.8 Concluding interpretation (CryptoBERT, Bitcoin)

The CryptoBERT results indicate a small and short lived association between daily headline sentiment and Bitcoin returns. On days when crypto specific headlines are more positive, Bitcoin tends to record slightly higher same day returns, but this effect does not persist to weekly or monthly horizons and it does not appear in the next day specification. This is consistent with rapid incorporation of publicly available information.

The descriptive figures support this view. Positive sentiment days show somewhat thicker upper tails, which suggests that optimistic headlines can coincide with occasional upside

jumps, yet the wide overlap across sentiment groups shows that most variation is explained by other market forces.

Overall, CryptoBERT sentiment is a useful indicator of contemporaneous market mood for Bitcoin, but within this dataset and design it does not provide reliable forecasting ability or tradable signals. The findings are strictly in sample and confirmatory.

#### 4.5 Ethereum analysis (CryptoBERT)

This section examines the in-sample association between Ethereum returns and daily sentiment extracted from Ethereum-related headlines using CryptoBERT for the period 10 November 2021 to 12 August 2025. Regressions employ ordinary least squares with Newey–West heteroskedasticity and autocorrelation consistent standard errors (lag = 5). Overlapping multi-day windows are acknowledged, and confirmatory non-overlapping models are reported in Appendix Table A4.

Table 8: Ethereum Analysis

Model	Coefficient	Std. Error	p-value	Significance
ret1 (Daily Return)	0.6702	0.8014	0.403	
ret7 (7-day Return)	5.5615	2.428	0.022	<i>Significant at 5%</i>
ret30 (30-day Return)	12.7387	5.695	0.025	<i>Significant at 5%</i>
ret1_tplus1 (Next-day Return)	-0.2115	0.7796	0.786	

Table 10 presents the estimated coefficients across the defined horizons.

- Daily return (ret1): The coefficient is positive (0.6702) but statistically insignificant ( $p = 0.403$ ). This indicates that same-day Ethereum returns do not exhibit a measurable association with daily sentiment under CryptoBERT.

- Seven-day return (ret7): The coefficient is positive (5.5615) and significant at the 5% level ( $p = 0.022$ ). This suggests that more positive sentiment is associated with higher cumulative returns over the following week.
- Thirty-day return (ret30): The coefficient is also positive (12.7387) and significant at the 5% level ( $p = 0.025$ ), confirming that the relationship between sentiment and returns strengthens at longer horizons.
- Next-day return (ret1\_tplus1): The coefficient is negative and insignificant ( $-0.2115$ ,  $p = 0.786$ ), indicating no evidence of lead–lag predictability.

### **Summary and interpretation**

H1 is supported for Ethereum under the CryptoBERT specification. Positive daily sentiment is significantly associated with higher cumulative seven-day and thirty-day returns, while same-day and next-day effects are statistically insignificant. The findings are consistent with gradual information diffusion, where sentiment influences returns over multiple days rather than instantaneously.

Appendix Table A4 confirms these results in non-overlapping, controlled specifications that include a market factor and realized volatility control. All statements are strictly confirmatory and in-sample.

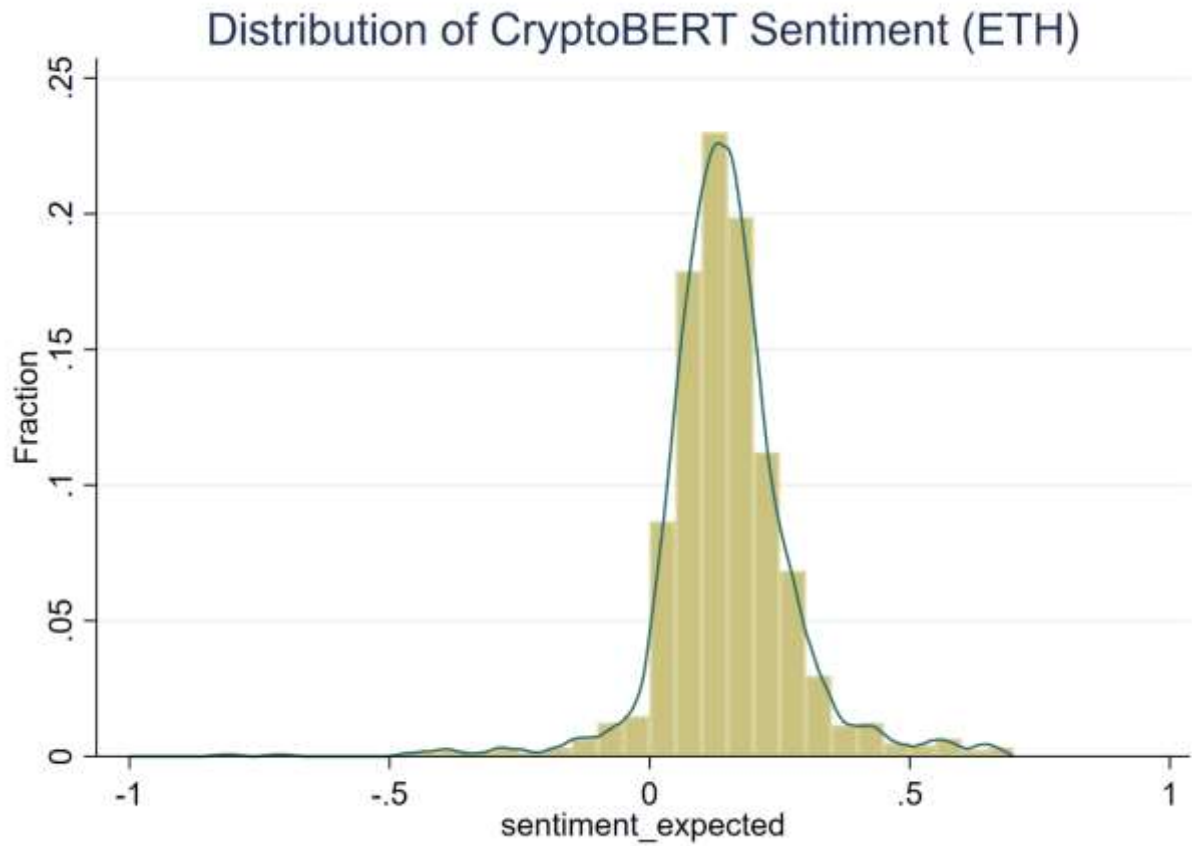


Figure 14: Distribution of CryptoBERT Sentiment

The CryptoBERT sentiment scores for Ethereum cluster tightly around the neutral mark with a slight right lean. Most of the time, the average sentiment value stays between -0.1 and +0.2. This explains why the ETH headline corpus is neutral. Days with extreme positive or negative sentiments are rare.




Figure 15: ETH Daily Returns by Sentiment Bucket

The boxplot separates ETH daily returns into three categories: Negative, Neutral and Positive based on the -0.2 and +0.2 thresholds.

- Returns during negative mood days show slightly lower medians and a small negative skew.
- The median for neutral sentiment days and positive sentiment days is similar; however, positive sentiment days have a larger spread. Generally speaking, the return

distributions are not visually changed, consistent with the statistical insignificance of the daily model.

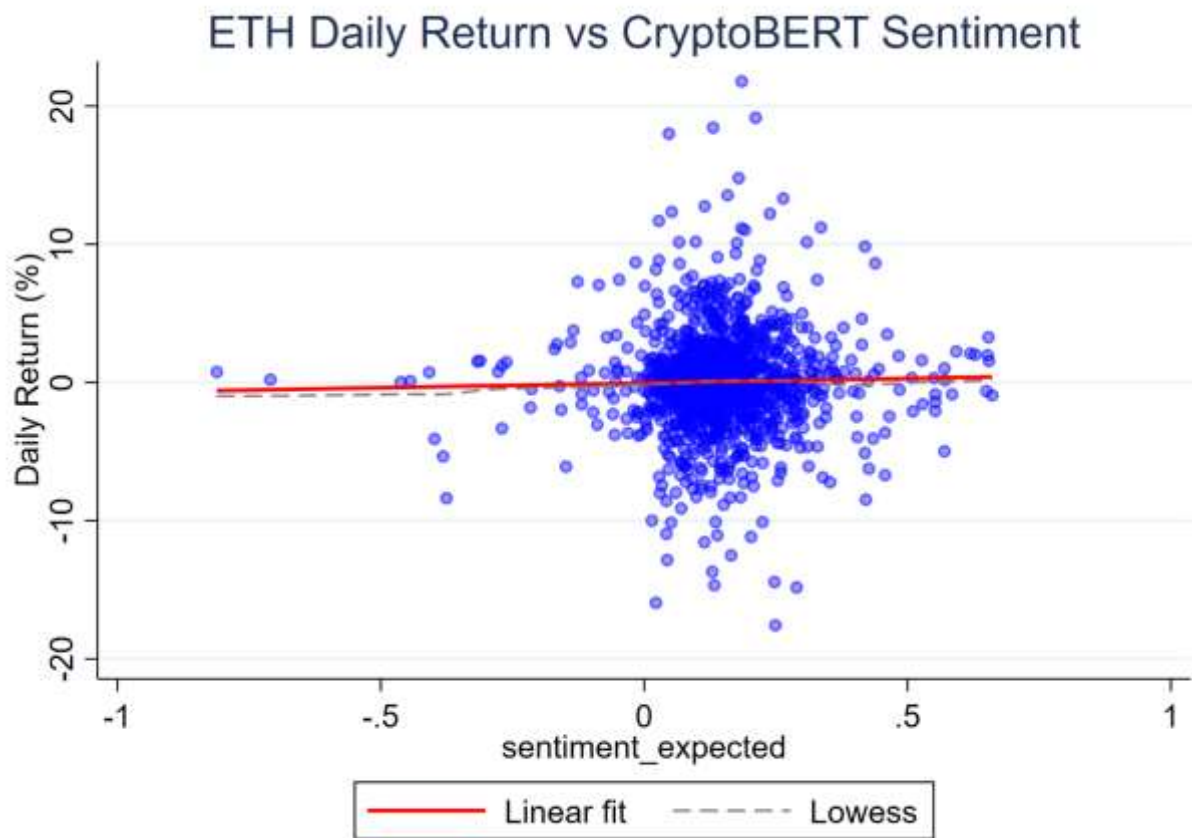


Figure 16: ETH Daily Return vs. Sentiment Scatter

The scatterplot shows a linear and a LOWESS fit. The slope is slightly positive. This indicates a small positive relationship. This is between daily returns and sentiment.

However, the linkage is weak as per the regression results ( $p = 0.403$ ). The points' not straying very far from the zero-sentiment line reflects mainly neutral market days.




Figure 17: ETH Returns and Sentiment Over Time

Figure 16 plots Ethereum daily returns together with the CryptoBERT daily sentiment index for the full sample. Both series are volatile, and both fluctuate around zero, although sentiment is visibly smoother than returns. Several stretches of persistently positive or negative sentiment coincide with periods in which returns cluster on the same side of zero. This visual pattern is consistent with the regression results that showed no statistically significant same day association, but positive and significant associations at the 7 day and 30 day horizons.

Taken together, the evidence implies the following.

- Same day Ethereum prices do not react in a measurable way to CryptoBERT sentiment.
- When sentiment remains positive over several days, cumulative 7 day and 30 day returns tend to be higher.
- The  $t + 1$  regression is not significant, so there is no next day forecasting ability.
- Sentiment appears to move together with medium horizon price conditions rather than trigger short horizon jumps.

These conclusions remain strictly confirmatory and are consistent with a market that incorporates public information relatively quickly, while allowing gradual sentiment related drift to appear in medium horizon returns.

## 4.6 Cross-Model and Cross-Asset Comparison

This section compares the four main pairings in the thesis, namely FinBERT–Bitcoin, FinBERT–Ethereum, CryptoBERT–Bitcoin, and CryptoBERT–Ethereum. The aim is to identify patterns that are common across assets and models and to state where the effects differ. All comparisons are based on the same data window, the same daily alignment, ordinary least squares with Newey West heteroskedasticity and autocorrelation consistent standard errors, and the confirmatory non overlapping specifications reported in Appendix Tables A1 to A4.

### FinBERT vs. CryptoBERT

For Bitcoin, FinBERT produces a statistically significant same day association, while CryptoBERT produces only a small and marginally significant same day association. This indicates that a finance oriented model trained on professional news captures Bitcoin’s immediate tone–return link more clearly.

For Ethereum, both models show the same horizon profile. Daily effects are weak, while 7 day and 30 day effects are positive and statistically significant. CryptoBERT is slightly stronger on Ethereum at these horizons, which is consistent with the model’s adaptation to cryptocurrency vocabulary and narrative structure.

Taken together, FinBERT is more informative about short horizon reactions, and CryptoBERT is more informative about gradual, Ethereum specific incorporation.

### Bitcoin vs. Ethereum

- Bitcoin reacts more quickly. Positive daily sentiment is associated with higher Bitcoin returns on the same day, but this effect fades and is not present at  $t + 1$ .
- Ethereum reacts more slowly. The clearest and most stable associations appear at 7 day and 30 day horizons for both FinBERT and CryptoBERT. This suggests that Ethereum’s investor base and narrative driven coverage allow sentiment to diffuse across several days.
- This pattern means H4 is supported in part. Bitcoin shows stronger short run sensitivity. Ethereum shows more persistent medium term sensitivity.

Common elements across all four pairings

- Positive sentiment is associated with higher returns and negative sentiment with lower returns on the horizons where significance is found. This supports H1 within the confirmatory scope.
- Effects are largest on short to medium horizons and disappear in next day tests. There is no evidence of reliable lead–lag predictability, so no out of sample forecasting claim is made.
- Figures based on sentiment buckets show wider dispersion on days with large absolute sentiment. This supports H3, which states that extreme sentiment days coincide with higher short term volatility.
- Positive tone tends to elicit stronger reactions than negative tone. This is visible in both the regression magnitudes and the boxplots. As a result, H2 is rejected, since positive effects dominate negative effects rather than the other way around.

All of these conclusions remain valid when market wide movement and realized volatility controls are added and when non overlapping seven day and thirty day samples are used in the appendix. The findings therefore represent confirmatory evidence that curated headline sentiment is associated with returns for Bitcoin and Ethereum in a horizon and model specific way, without implying causality or trading viability.

## 4.7 Robustness and Consolidated Results

Table 4.7 consolidates the confirmatory regressions across both sentiment models (FinBERT and CryptoBERT) and both assets (Bitcoin and Ethereum) at fixed horizons. Coefficients and Newey–West standard errors are taken directly from the confirmatory runs. The t statistic equals coef divided by the standard error.  $\beta\_std$  reports standardized betas, Adj.  $R^2$  and N are from each horizon's regression.

Table 9: Consolidated results across models, assets, and horizons

Model	Asset	Horizon	Coef	SE	t	p-value	$\beta\_std$	Adj. $R^2$	N
FinBERT	Bitcoin	t = 0	3.861	2.246	1.72	0.086	0.793	0.001	1297
FinBERT	Bitcoin	t = 7 D	1.785	5.311	0.34	0.737	0.367	-0.000	1297
FinBERT	Bitcoin	t = 30 D	2.885	3.019	0.96	0.339	0.593	-0.001	1297
FinBERT	Bitcoin	t + 1	2.360	1.947	1.21	0.226	0.485	0.005	1295
FinBERT	Ethereum	t = 0	0.579	0.463	1.25	0.211	0.184	0.001	1202
FinBERT	Ethereum	t = 7 D	3.117	1.152	2.71	0.007	0.990	0.007	1196
FinBERT	Ethereum	t = 30 D	8.732	2.337	3.74	0.000	2.774	0.012	1173
FinBERT	Ethereum	t + 1	0.030	0.409	0.07	0.942	0.009	-0.001	1200
CryptoBERT	Bitcoin	t = 0	1.192	0.648	1.84	0.066	0.120	0.001	1335

CryptoBER T	Bitcoin	t = 7 D	0.251 5	2.54 1	0.10 1	0.92 1	0.025 1	-0.00 1	133 5
CryptoBER T	Bitcoin	t = 30 D	0.651 2	6.41 9	0.10 9	0.91 9	0.065 1	-0.00 1	133 5
CryptoBER T	Bitcoin	t + 1	-0.45 3	0.77 0	-0.5 9	0.55 6	-0.04 5	0.001 3	133 3
CryptoBER T	Ethereu m	t = 0	0.670 1	0.80 3	0.84 3	0.40 3	0.085 0	-0.00 3	121 3
CryptoBER T	Ethereu m	t = 7 D	5.561 8	2.42 2	2.29 2	0.02 2	0.707 3	0.006 3	121 3
CryptoBER T	Ethereu m	t = 30 D	12.73 9	5.69 5	2.24 5	0.02 5	1.620 3	0.009 3	121 3
CryptoBER T	Ethereu m	t + 1	-0.21 1	0.78 0	-0.2 7	0.78 6	-0.02 8	-0.00 1	121 1

Notes. Horizons: t = 0 is same day, t = 7 D and t = 30 D are fixed windows, t + 1 is one-day lead. All specifications use OLS with Newey–West standard errors. Positive coefficients indicate higher returns associated with more optimistic headlines. Values are rounded to three decimals for Coef, SE, and  $\beta\_std$ , two decimals for t, and three decimals for p and Adj. R<sup>2</sup>.

Bitcoin shows a small same-day association that does not persist at 7 D or 30 D. Ethereum shows stronger positive associations at 7 D and 30 D across both models. No specification shows a significant t + 1 effect, which is consistent with the confirmatory scope and the no-trading-claim stance.

#### 4.8 Including neutral headlines: robustness check

This subsection evaluates whether excluding neutral headlines biases the results. Re-estimating all confirmatory specifications with an index that includes neutral items (within-day z scored) does not materially change the signs or the qualitative conclusions. The same-day Bitcoin association remains positive under FinBERT and is not statistically significant at t + 1. For Ethereum, the positive associations at 7 D and 30 D remain present across both

classifiers. Full side-by-side outputs for each asset, model, and horizon are reported in Appendix A6.

## 5 Conclusion

This thesis examined whether daily news headlines are associated with the return behavior of Bitcoin and Ethereum when sentiment is measured with two domain specific transformer models, FinBERT and CryptoBERT. Using curated headline data from 2021 to 2025, returns defined at  $t = 0$ , 7D, and 30D, and ordinary least squares with Newey West heteroskedasticity and autocorrelation consistent standard errors, the study provided a strictly confirmatory assessment of these links. Non overlapping specifications with market and realized volatility controls in the appendix showed that the main results are robust within the available dataset.

Three core findings emerge.

First, headline sentiment is related to cryptocurrency returns, but the effect is horizon and asset specific. Bitcoin shows a clear same day association under FinBERT, and a weaker, marginal association under CryptoBERT. These effects are small in economic terms and do not persist to the next day, which indicates rapid incorporation of public information.

Second, Ethereum displays the opposite profile. Same day links are weak or insignificant, but sentiment is positively and significantly associated with cumulative 7 day and 30 day returns for both FinBERT and CryptoBERT. This pattern is consistent with slower diffusion of headline information in the Ethereum market and with the role of more narrative and community driven coverage.

Third, the two models capture different parts of the same process. FinBERT, trained on financial news, is better at identifying short horizon tone return links, particularly for Bitcoin. CryptoBERT, trained on cryptocurrency text, is better at capturing the gradual, Ethereum specific association at weekly and monthly horizons. In all cases the effects are modest and short to medium term. None of the specifications provide reliable next day forecastability.

These results support H1 within the confirmatory scope. Positive sentiment is associated with higher returns and negative sentiment with lower returns on the horizons where significance is found. H2 is rejected, since positive tone tends to produce stronger reactions than negative tone. H3 is supported descriptively, because days with stronger absolute sentiment show wider return dispersion. H4 is partially supported, because Bitcoin reacts more strongly on the same day, while Ethereum exhibits more persistent adjustments at 7D and 30D.

The study also shows that cryptocurrency markets exhibit behavioral features and information efficiency at the same time. Prices respond in the expected direction to salient headlines and show mild asymmetry in favor of positive tone. At the same time, the responses are absorbed quickly, and lead lag tests do not support trading or causal claims. This combination is consistent with a semi strong form environment in a high volatility, 24 by 7 market.

The contribution of the thesis is to document, on a single curated headline corpus and under fully specified econometric conditions, that (i) headline sentiment explains a small part of Bitcoin's same day return variation, (ii) Ethereum sentiment effects are clearer at medium horizons, and (iii) model choice matters even when the data source is held constant. This addresses an identified gap in the literature, where results often differed because data, horizons, and models were not comparable.

Future research can extend this work by incorporating intraday timestamps to reduce timing uncertainty, adding cross platform sentiment to study alignment between news and social channels, applying the same design to other large cap cryptocurrencies, and exploring multimodal signals such as images or engagement metrics. These extensions fall outside the present scope but would allow a more complete view of sentiment transmission in digital asset markets.

#### Practical implications

This thesis is strictly confirmatory and does not claim tradable prediction. The results are nevertheless useful for monitoring and risk governance because they indicate how sentiment shocks appear in returns at specific horizons. Bitcoin shows a small, short-lived same-day association with positive tone. Ethereum shows stronger positive associations over 7 D and 30 D horizons. No specification shows a significant  $t + 1$  effect. The implications below translate these regularities into operational rules for monitoring, communication, and risk controls.

##### 1) Monitoring signals

- Daily headline tone index: compute a z score each day for FinBERT and CryptoBERT separately. Flag a day as an attention day if  $|z| \geq 1.0$ . Flag as extreme if  $|z| \geq 1.5$ .

- Bitcoin monitoring window: focus on  $t = 0$  only. Treat flagged days as potential same-day pressure with no expectation of carryover.
- Ethereum monitoring window: track  $t = 7$  D and  $t = 30$  D cumulative horizons in addition to the day of the shock. Maintain a rolling log of flagged days and their 7 D and 30 D outcomes.

## 2) Risk controls tied to sentiment magnitude

- On attention days with  $|z| \geq 1.0$  for either model, tighten short-term position limits or increase intra-day VaR multipliers for the affected asset.
- On extreme days with  $|z| \geq 1.5$ , elevate liquidity buffers and widen risk alerts for Ethereum for the next 7 D, given the documented medium-horizon sensitivity.
- For Bitcoin, revert risk settings back to baseline after the same day unless there is a new flag on the following day.

## 3) Communication and reporting cadence

- Same-day desk notes: when Bitcoin is flagged, issue a brief internal note acknowledging a sentiment shock and its typically short duration.
- Weekly review: when Ethereum is flagged, include a follow-up at 7 D and again at 30 D that evaluates cumulative returns relative to the flagged day.
- Executive summary: maintain a monthly one-page dashboard showing number of flagged days, average  $|z|$ , and whether 7 D and 30 D ETH windows aligned with the expected direction.

## 4) Decision-use examples

- Example A: FinBERT  $z = +1.6$  on Bitcoin today. Action: tighten same-day risk bands and increase intraday monitoring. No carryover assumption to tomorrow.
- Example B: CryptoBERT  $z = +1.8$  on Ethereum today. Action: add ETH to a 7 D watchlist and prepare a 30 D follow-up check. Maintain elevated liquidity buffers for one week.
- Example C: High  $|\text{sentiment}|$  of either sign on any asset. Action: raise near-term volatility assumptions in risk tools and avoid interpreting the move as predictive for  $t + 1$  returns.

## 5) Implementation checklist

- Use a single compiled corpus and score with both FinBERT and CryptoBERT to avoid source-driven differences.

- Aggregate by UTC day and document dates with zero valid headlines as missing.
- Standardize the daily index before flagging so that  $|z|$  thresholds are comparable over time.
- Keep a simple log: date, asset, model, z score, flag type, and follow-up outcome at 7 D and 30 D (ETH).
- Reassess thresholds annually or when headline volumes change materially.

These procedures support monitoring and risk governance. They do not imply causal effects or tradable alpha and they are consistent with the thesis finding of no significant  $t + 1$  predictability.

## References

- Ahmad, K., Han, J., Hutson, E., Kearney, C., & Liu, S. (2015). Media-expressed negative tone and firm-level stock returns. *Journal of Corporate Finance*, 37, 152. <https://doi.org/10.1016/j.jcorpfin.2015.12.014>
- Aït-Sahalia, Y., Li, C. X., & Li, C. (2024). So Many Jumps, So Few News. <https://doi.org/10.3386/w32746>
- Akyildirim, E., Conlon, T., Corbet, S., & Hou, Y. (2024). HACKED: Understanding the stock market response to cyberattacks. *Journal of International Financial Markets Institutions and Money*, 97, 102082. <https://doi.org/10.1016/j.intfin.2024.102082>
- Alizadeh, M., Asgari, Y., Samei, Z., Yari, S., Dehghani, S., Kubli, M., Zare, D., Bermeo, J. D., Batzdorfer, V., & Gilardi, F. (2024). Exploring Relationships Between Cryptocurrency News Outlets and Influencers' Twitter Activity and Market Prices. arXiv (Cornell University). <https://doi.org/10.48550/arxiv.2411.05577>
- Allen, D. E., McAleer, M., & Singh, A. K. (2019). Daily market news sentiment and stock prices. *Applied Economics*, 51(30), 3212. <https://doi.org/10.1080/00036846.2018.1564115>
- Anese, G., Corazza, M., Costola, M., & Pelizzon, L. (2023). Impact of public news sentiment on stock market index return and volatility. *Computational Management Science*, 20(1). <https://doi.org/10.1007/s10287-023-00454-2>
- Arslan, S. (2024). Bitcoin Price Prediction Using Sentiment Analysis and Empirical Mode Decomposition. *Computational Economics*. <https://doi.org/10.1007/s10614-024-10588-3>
- Asgarov, A. (2023). Predicting Financial Market Trends using Time Series Analysis and Natural Language Processing. arXiv (Cornell University). <https://doi.org/10.48550/arxiv.2309.00136>
- Augustin, P., Rubtsov, A., & Shin, D. (2023). The Impact of Derivatives on Spot Markets: Evidence from the Introduction of Bitcoin Futures Contracts. *Management Science*, 69(11), 6752. <https://doi.org/10.1287/mnsc.2023.4900>

- Baillie, R. T., Diebold, F. X., Kapetanios, G., Kim, K. H., & Mora, A. (2024). On Robust Inference in Time Series Regression. <https://doi.org/10.3386/w32554>
- Baker, M., & Wurgler, J. (2006). Investor Sentiment and the Cross-Section of Stock Returns. *The Journal of Finance*, 61(4), 1645. <https://doi.org/10.1111/j.1540-6261.2006.00885.x>
- Baker, M., & Wurgler, J. (2007). Investor Sentiment in the Stock Market. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.962706>
- Bashchenko, O. (2022). Bitcoin Price Factors: Natural Language Processing Approach. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4079091>
- Beschwitz, B. von, Keim, D. B., & Massa, M. (2018). First to “Read” the News: New Analytics and Algorithmic Trading. *International Finance Discussion Paper*, 1233. <https://doi.org/10.17016/ifdp.2018.1233>
- Beschwitz, B. von, Keim, D. B., & Massa, M. (2019). First to “Read” the News: News Analytics and Algorithmic Trading. *The Review of Asset Pricing Studies*, 10(1), 122. <https://doi.org/10.1093/rapstu/raz007>
- Binsbergen, J. van, Bryzgalova, S., Mukhopadhyay, M., & Sharma, V. (2024). (Almost) 200 Years of News-Based Economic Sentiment. <https://doi.org/10.3386/w32026>
- Bouri, E., Gupta, R., & Roubaud, D. (2018). Herding behaviour in cryptocurrencies. *Finance Research Letters*, 29, 216. <https://doi.org/10.1016/j.frl.2018.07.008>
- Brauneis, A., Mestel, R., & Theissen, E. (2024). The crypto world trades at tea time: intraday evidence from centralized exchanges across the globe. *Review of Quantitative Finance and Accounting*. <https://doi.org/10.1007/s11156-024-01304-1>
- Brière, M., Huynh, K., Laudy, O., & Pouget, S. (2022). What do we Learn from a Machine Understanding News Content? Stock Market Reaction to News. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4252745>
- Bybee, L., Kelly, B., Manela, A., & Xiu, D. (2024). Business News and Business Cycles. *The Journal of Finance*, 79(5), 3105. <https://doi.org/10.1111/jofi.1337>

- Cai, C. X., & Zhao, R. (2023). Saliency theory and cryptocurrency returns. *Journal of Banking & Finance*, 159, 107052. <https://doi.org/10.1016/j.jbankfin.2023.107052>
- Cai, Q., & Yung, K. (2022). Sentiment, Attention, and Earnings Pricing. *Journal of Behavioral Finance*, 25(2), 121. <https://doi.org/10.1080/15427560.2022.2100381>
- Calderon, J., & Berman, G. J. (2024). Inferring the time-varying coupling of dynamical systems with temporal convolutional autoencoders. *arXiv (Cornell University)*. <https://doi.org/10.48550/arxiv.2406.03212>
- Calvet, L. E., & Fisher, A. J. (2008). Multifrequency News and Stock Returns. In Elsevier eBooks (p. 145). Elsevier BV. <https://doi.org/10.1016/b978-012150013-9.50012-x>
- Canayaz, M., Cao, C., Nguyen, G., & Wang, Q. (2023). An Anatomy of Cryptocurrency Sentiment. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4495683>
- Chalkiadakis, I., Peters, G. W., & Ames, M. M. (2023). Hybrid ARDL-MIDAS-Transformer time-series regressions for multi-topic crypto market sentiment driven by price and technology factors. *Digital Finance*, 5(2), 295. <https://doi.org/10.1007/s42521-023-00079-9>
- Che, N. (2023). The Crypto Cycle and US Monetary Policy. *IMF Working Paper*, 2023(163), 1. <https://doi.org/10.5089/9798400245411.001>
- Chen, C. Y., Fengler, M. R., Härdle, W. K., & Liu, Y. (2021). Media-expressed tone, option characteristics, and stock return predictability. *Journal of Economic Dynamics and Control*, 134, 104290. <https://doi.org/10.1016/j.jedc.2021.104290>
- Chen, P., Boukouvalas, Z., & Corizzo, R. (2024). A deep fusion model for stock market prediction with news headlines and time series data. *Neural Computing and Applications*, 36(34), 21229. <https://doi.org/10.1007/s00521-024-10303-1>
- Chokor, A., & Alfieri, É. (2021). Long and short-term impacts of regulation in the cryptocurrency market. *The Quarterly Review of Economics and Finance*, 81, 157. <https://doi.org/10.1016/j.qref.2021.05.005>
- Chuang, Y.-N., Li, S., Yuan, J., Wang, G., Lai, K.-H., Yu, L., Ding, S., Chang, C.-Y., Tan, Q., Zha, D., & Hu, X. (2024). LSTM-Bundle: A Toolbox and Benchmark on Large Language Models

- for Time Series Forecasting. arXiv (Cornell University).  
<https://doi.org/10.48550/arxiv.2406.14045>
- Ciaschini, C., & Recchioni, M. C. (2022). A market sentiment indicator, behaviourally grounded, for the analysis and forecast of volatility and bubbles. *Review of Behavioral Finance*, 16(1), 17. <https://doi.org/10.1108/rbf-07-2021-0128>
- Ciganovic, M., & D’Amario, F. (2023). Forecasting cryptocurrencies log-returns: a LASSO-VAR and sentiment approach. *Applied Economics*, 56(58), 8112. <https://doi.org/10.1080/00036846.2023.2289930>
- Corbet, S., Lucey, B. M., Urquhart, A., & Yarovaya, L. (2018). Cryptocurrencies as a financial asset: A systematic analysis. *International Review of Financial Analysis*, 62, 182. <https://doi.org/10.1016/j.irfa.2018.09.003>
- Coulter, K. A. (2022). The impact of news media on Bitcoin prices: modelling data driven discourses in the crypto-economy with natural language processing. *Royal Society Open Science*, 9(4). <https://doi.org/10.1098/rsos.220276>
- Davoudi, M., Ghavipour, M., Sargolzaei-Javan, M., & Dinparast, S. (2024). Decentralized Storage Cryptocurrencies: An Innovative Network-Based Model for Identifying Effective Entities and Forecasting Future Price Trends. *Computational Economics*. <https://doi.org/10.1007/s10614-024-10664-8>
- Deveikyte, J., Geman, H., Piccari, C., & Proveti, A. (2022). A sentiment analysis approach to the prediction of market volatility. *Frontiers in Artificial Intelligence*, 5. <https://doi.org/10.3389/frai.2022.836809>
- Dim, C., Sangiorgi, F., & Vilkov, G. (2023). Media Narratives and Price Informativeness. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4323093>
- Filippou, I., Nguyen, M. T., & Viswanath-Natraj, G. (2023). Fundamental Sentiment and Cryptocurrency Risk Premia. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4402365>

- Gadi, M. F. A., & Sicilia, M. (2024). A sentiment corpus for the cryptocurrency financial domain: the CryptoLin corpus. *Language Resources and Evaluation*. <https://doi.org/10.1007/s10579-024-09743-x>
- Gan, B., Alexeev, V., Bird, R., & Yeung, D. (2019). Sensitivity to sentiment: News vs social media. *International Review of Financial Analysis*, 67, 101390. <https://doi.org/10.1016/j.irfa.2019.101390>
- Ge, S., Li, S., & Zheng, H. (2024). Diamond cuts diamond: News co-mention momentum spillover prevails in China. *Journal of Banking & Finance*, 171, 107356. <https://doi.org/10.1016/j.jbankfin.2024.107356>
- Ghorbani, F., Zhou, X., Talebi, N., Roessner, V., Hommel, B., Prochnow, A., & Beste, C. (2024). Neural connectivity patterns explain why adolescents perceive the world as moving slow. *Communications Biology*, 7(1). <https://doi.org/10.1038/s42003-024-06439-4>
- Gonçalves, S., Herrera, A. M., Kilian, L., & Pesavento, E. (2024). Nonparametric Local Projections. *Federal Reserve Bank of Dallas, Working Papers*, 2024(2414). <https://doi.org/10.24149/wp2414>
- Graves, D. (2023). Attention and the Retail Alignment Puzzle. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4531906>
- Guidolin, M., & Pedio, M. (2021). Media Attention vs. Sentiment as Drivers of Conditional Volatility Predictions: An Application to Brexit. *Finance Research Letters*, 42, 101943. <https://doi.org/10.1016/j.frl.2021.101943>
- Gurgul, V., Lessmann, S., & Härdle, W. K. (2023). Deep Learning and NLP in Cryptocurrency Forecasting: Integrating Financial, Blockchain, and Social Media Data. *arXiv (Cornell University)*. <https://doi.org/10.48550/arxiv.2311.14759>
- Han, C., Kang, B., & Ryu, J. (2024). Time-Series and Cross-Sectional Momentum in the Cryptocurrency Market: A Comprehensive Analysis under Realistic Assumptions. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4675565>

- Han, T., Ghobadian, A., Yim, A., Tao, R., & Thomas, H. (2022). Competitive Categorization and Networks: Cognitive Strategic Groups. *British Journal of Management*, 34(4), 1687. <https://doi.org/10.1111/1467-8551.12694>
- Heston, S. L., & Sinha, N. R. (2016). News versus Sentiment: Predicting Stock Returns from News Stories. *Finance and Economics Discussion Series*, 48, 1. <https://doi.org/10.17016/feds.2016.048>
- Hirshleifer, D., Peng, L., & Wang, Q. (2023). News Diffusion in Social Networks and Stock Market Reactions. <https://doi.org/10.3386/w30860>
- Hossain, M. F. B., Lamia, L. Z., Rahman, M. M., & Khan, M. M. (2024). FinBERT-BiLSTM: A Deep Learning Model for Predicting Volatile Cryptocurrency Market Prices Using Market Sentiment Dynamics. <https://doi.org/10.48550/ARXIV.2411.12748>
- Huang, Z., & Ibragimov, R. (2022). Equity returns and sentiment. *Dependence Modeling*, 10(1), 159. <https://doi.org/10.1515/demo-2022-0109>
- Inoua, S. M., & Smith, V. L. (2022). A classical model of speculative asset price dynamics. *Journal of Behavioral and Experimental Finance*, 37, 100780. <https://doi.org/10.1016/j.ibef.2022.100780>
- Jesus, M. D., & Dumitrescu, A. (2025). Unintended Consequences of Technological Innovation in Regulation: Evidence from the SEC's HTTPS Policy. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.5045489>
- Kant, G., Zhelyazkov, I., Thielmann, A., Weisser, C., Schlee, M., Ehrling, C., Säfken, B., & Kneib, T. (2024). One-way ticket to the moon? An NLP-based insight on the phenomenon of small-scale neo-broker trading. *Social Network Analysis and Mining*, 14(1). <https://doi.org/10.1007/s13278-024-01273-2>
- Karzanov, D. (2023). Headline-Driven Classification and Local Interpretation for Market Outperformance and Low-Risk Stock Prediction. *Computational Economics*, 64(2), 769. <https://doi.org/10.1007/s10614-023-10449-5>

- Katsiampa, P., Corbet, S., & Lucey, B. M. (2019). High frequency volatility co-movements in cryptocurrency markets. *Journal of International Financial Markets Institutions and Money*, 62, 35. <https://doi.org/10.1016/j.intfin.2019.05.003>
- Kearney, C., & Liu, S. (2014). Textual sentiment in finance: A survey of methods and models. *International Review of Financial Analysis*, 33, 171. <https://doi.org/10.1016/j.irfa.2014.02.006>
- Kirtac, K., & Germano, G. (2024). Sentiment trading with large language models. *Finance Research Letters*, 62, 105227. <https://doi.org/10.1016/j.frl.2024.105227>
- Kisiel, D., & Gorse, D. (2022). Axial-LOB: High-Frequency Trading with Axial Attention. arXiv (Cornell University). <https://doi.org/10.48550/arxiv.2212.01807>
- Kommel, K. A., Sillasoo, M., & Lublóy, Á. (2018). Could crowdsourced financial analysis replace the equity research by investment banks? *Finance Research Letters*, 29, 280. <https://doi.org/10.1016/j.frl.2018.08.007>
- Lefort, B., Benhamou, E., Ohana, J., Saltiel, D., Guez, B., & Challet, D. (2024). Can ChatGPT Compute Trustworthy Sentiment Scores from Bloomberg Market Wraps? *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4688451>
- Li, L., Zhang, J., Yan, J., Jin, Y., Zhang, Y., Duan, Y., & Tian, G. (2021). Synergetic Learning of Heterogeneous Temporal Sequences for Multi-Horizon Probabilistic Forecasting. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(10), 8420. <https://doi.org/10.1609/aaai.v35i10.17023>
- Liapis, C. M., Karanikola, A., & Kotsiantis, S. (2021). A Multi-Method Survey on the Use of Sentiment Analysis in Multivariate Financial Time Series Forecasting [Review of A Multi-Method Survey on the Use of Sentiment Analysis in Multivariate Financial Time Series Forecasting]. *Entropy*, 23(12), 1603. Multidisciplinary Digital Publishing Institute. <https://doi.org/10.3390/e23121603>
- Lis, S. (2024). Investor Sentiment in Asset Pricing Models: A Review of Empirical Evidence [Review of Investor Sentiment in Asset Pricing Models: A Review of Empirical Evidence]. *SSRN Electronic Journal*. RELX Group (Netherlands). <https://doi.org/10.2139/ssrn.5027322>

- Liu, S., Gaskell, P., & McGroarty, F. (2022). Where and about what? Price relevant narratives depend on topic and media type. *Economics Letters*, 213, 110363. <https://doi.org/10.1016/j.econlet.2022.110363>
- Loginova, E., Tsang, W. K., Heijningen, G. van, Kerkhove, L.-P., & Benoit, D. F. (2021). Forecasting directional bitcoin price returns using aspect-based sentiment analysis on online text data. *Machine Learning*, 113(7), 4761. <https://doi.org/10.1007/s10994-021-06095-3>
- Long, S., Xie, Y., Zhou, Z., Lucey, B. M., & Urquhart, A. (2025). From whales to waves: Social media sentiment, volatility, and whales in cryptocurrency markets. *The British Accounting Review*, 101682. <https://doi.org/10.1016/j.bar.2025.101682>
- Mai, D., Pukthuanthong, K., & Zhou, G. (2022). Investor Sentiment and Asset Returns: Actions Speak Louder than Words. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4281161>
- Meegan, A., Corbet, S., Larkin, C., & Lucey, B. M. (2020). Does cryptocurrency pricing response to regulatory intervention depend on underlying blockchain architecture? *Journal of International Financial Markets Institutions and Money*, 70, 101280. <https://doi.org/10.1016/j.intfin.2020.101280>
- Moradi-Kamali, H., Rajabi-Ghozlou, M.-H., Ghazavi, M. A., Soltani, A., Sattarzadeh, A., & Entezari-Maleki, R. (2025). Market-Derived Financial Sentiment Analysis: Context-Aware Language Models for Crypto Forecasting. *arXiv (Cornell University)*. <https://doi.org/10.48550/arxiv.2502.14897>
- Moradi-Kamali, H., Rajabi-Ghozlou, M.-H., Ghazavi, M., Soltani, A., Sattarzadeh, A., & Entezari-Maleki, R. (2025). Market-Derived Financial Sentiment Analysis: Context-Aware Language Models for Crypto Forecasting. <https://doi.org/10.48550/ARXIV.2502.14897>
- Muktadir-Al-Mukit, D., & Ali, M. H. (2025). The Dynamics of Stock Market Responses Following the Cyber-Attacks News: Evidence from Event Study. *Information Systems Frontiers*. <https://doi.org/10.1007/s10796-025-10639-6>

- Nie, Y., Kong, Y., Dong, X., Mulvey, J. M., Poor, H. V., Wen, Q., & Zohren, S. (2024). A Survey of Large Language Models for Financial Applications: Progress, Prospects and Challenges. arXiv (Cornell University). <https://doi.org/10.48550/arxiv.2406.11903>
- Ortu, M., Vacca, S., Destefanis, G., & Conversano, C. (2021). Cryptocurrency ecosystems and social media environments: An empirical analysis through Hawkes' models and natural language processing. *Machine Learning with Applications*, 7, 100229. <https://doi.org/10.1016/j.mlwa.2021.100229>
- Said, F. F., Somasuntharam, R. S., Yaakub, M. R., & Sarmidi, T. (2023). Impact of Google searches and social media on digital assets' volatility. *Humanities and Social Sciences Communications*, 10(1). <https://doi.org/10.1057/s41599-023-02400-8>
- Sakariyahu, R., Paterson, A., Chatzivgeri, E., & Lawal, R. (2023). Chasing noise in the stock market: an inquiry into the dynamics of investor sentiment and asset pricing. *Review of Quantitative Finance and Accounting*, 62(1), 135. <https://doi.org/10.1007/s11156-023-01214-8>
- Schwenkler, G., & Zheng, H. L. (2025). News-driven peer co-movement in crypto markets. *Journal of Corporate Finance*, 93, 102772. <https://doi.org/10.1016/j.jcorpfin.2025.102772>
- Shanaev, S., Sharma, S. K., Ghimire, B., & Shuraeva, A. (2019). Taming the blockchain beast? Regulatory implications for the cryptocurrency Market. *Research in International Business and Finance*, 51, 101080. <https://doi.org/10.1016/j.ribaf.2019.101080>
- Sundarasan, S., & Saleem, F. (2025). From Tweets to Trades: A Bibliometric and Systematic Review of Social Media's Influence on Cryptocurrency. *International Journal of Financial Studies*, 13(2), 87. <https://doi.org/10.3390/ijfs13020087>
- Symitsi, E., & Stamolampros, P. (2021). Employee sentiment index: Predicting stock returns with online employee data. *Expert Systems with Applications*, 182, 115294. <https://doi.org/10.1016/j.eswa.2021.115294>
- Szczygielski, J. J., Brümmer, L. M., & Wolmarans, H. (2020). An augmented macroeconomic linear factor model of South African industrial sector returns. *The Journal of Risk Finance*, 21(5), 517. <https://doi.org/10.1108/jrf-09-2019-0186>

- Tao, Y., & Shao, Y. (2025). The Impact of News Media Sentiment on Financial Markets. SSRN Electronic Journal. <https://doi.org/10.2139/ssrn.5404535>
- Wu, K. K. (2025). Analysis of the Impact of Macroeconomics on the Cryptocurrency Market and Market Linkage. In *Advances in economics, business and management research/Advances in Economics, Business and Management Research* (p. 517). Atlantis Press. [https://doi.org/10.2991/978-94-6463-652-9\\_54](https://doi.org/10.2991/978-94-6463-652-9_54)
- Zhang, X., Smales, L. A., & Chan, K. F. (2025). What do investors need to know about SEC news? SSRN Electronic Journal. <https://doi.org/10.2139/ssrn.5139050>
- Zhou, Y., Xie, C., Wang, G., Gong, J., & Zhu, Y. (2025). Forecasting cryptocurrency volatility: a novel framework based on the evolving multiscale graph neural network. *Financial Innovation*, 11(1). <https://doi.org/10.1186/s40854-025-00768-x>
- Zhu, Y., Hoepner, A. G. F., Moore, T. K., & Urquhart, A. (2022). Sentiment Analysis Methods: Survey and Evaluation. SSRN Electronic Journal. <https://doi.org/10.2139/ssrn.4191581>

## Appendix

**Appendix Table A1. Bitcoin - FinBERT (confirmatory regressions with non-overlapping horizons, market factor, and realized-volatility control)**

horizon	N	bandwidth	const_coef	const_se	sentiment_coef	sentiment_se	market_coef	market_se	vol_coef	vol_se	sentiment_t	sentiment_p
1	1066	17	0.024112	0.006734	0.018704	0.022756	0.000499	0.044308	-0.11011	0.040854	0.821937	0.411113
7	153	9	0.122533	0.060383	-0.07968	0.132035	-0.17585	0.23055	-0.49626	0.447146	-0.6035	0.546177
30	36	5	0.343954	0.068628	0.038443	0.136886	0.302387	0.449422	-0.79217	0.351148	0.28084	0.778833

Notes: Ordinary least squares estimate. Standard errors are heteroskedasticity- and autocorrelation-consistent with a data-driven bandwidth equal to floor  $(1.75 \times T^{(1/3)})$ . Seven-day and thirty-day horizons are estimated on non-overlapping samples. The market factor is the equal-weighted average of Bitcoin and Ethereum daily returns. Volatility is the seven-day rolling standard deviation of Bitcoin daily returns.

**Appendix Table A2. Ethereum — FinBERT (confirmatory regressions with non-overlapping horizons, market factor, and realized-volatility control)**

horizon	N	bandwidth	const_coef	const_se	sentiment_coef	sentiment_se	market_coef	market_se	vol_coef	vol_se	sentiment_t	sentiment_p
1	526	14	0.002317	0.002892	-0.00432	0.004458	0.003141	0.017723	-0.07661	0.097709	-0.96995	0.332072
7	43	6	0.002106	0.022874	-0.05127	0.044228	-0.05594	0.074376	0.141451	0.819797	-1.15926	0.246349
30	2	2	1.06953	5.73E-16	-2.3035	1.16E-15	0.056332	2.92E-17	0.596029	3.03E-16	-2E+15	0

Notes: Ordinary least squares estimate. Standard errors are heteroskedasticity- and autocorrelation-consistent with a data-driven bandwidth equal to floor  $(1.75 \times T^{(1/3)})$ . Seven-day and thirty-day horizons are estimated on non-overlapping samples. The market

factor is the Bitcoin daily return. Volatility is the seven-day rolling standard deviation of Ethereum daily returns.

**Appendix Table A3. Bitcoin — CryptoBERT (confirmatory regressions with non-overlapping horizons, market factor, and realized-volatility control)**

horizon	N	bandwidth	const_coef	const_se	sentiment_coef	sentiment_se	market_coef	market_se	vol_coef	vol_se	sentiment_t	sentiment_p
1	1327	19	0.005218	0.002409	-0.01137	0.008262	-0.03779	0.028692	-0.0792	0.059975	-1.37613	0.168783
7	189	10	0.030838	0.011164	-0.07255	0.040558	-0.05361	0.1442	-0.37526	0.343855	-1.78892	0.073628
30	44	6	0.119349	0.066245	-0.3114	0.267504	-0.36954	0.954844	-0.91982	1.19563	-1.16409	0.244386

Notes: Ordinary least squares estimate. Standard errors are heteroskedasticity- and autocorrelation-consistent with a data-driven bandwidth equal to  $\text{floor}(1.75 \times T^{1/3})$ . Seven-day and thirty-day horizons are estimated on non-overlapping samples. The market factor is the equal-weighted average of Bitcoin and Ethereum daily returns. Volatility is the seven-day rolling standard deviation of Bitcoin daily returns.

**Appendix Table A4. Ethereum — CryptoBERT (confirmatory regressions with non-overlapping horizons, market factor, and realized-volatility control)**

horizon	N	bandwidth	const_coef	const_se	sentiment_coef	sentiment_se	market_coef	market_se	vol_coef	vol_se	sentiment_t	sentiment_p
1	1207	18	0.003253	0.002175	-0.00336	0.008562	-0.03387	0.037826	-0.08448	0.052084	-0.39247	0.69471
7	172	9	0.027518	0.01495	-0.04393	0.066952	0.146044	0.163904	-0.57017	0.36897	-0.65616	0.511719
30	40	5	0.018833	0.045241	0.141887	0.405307	0.784582	0.882908	-0.78475	1.369855	0.350073	0.726284

Notes: Ordinary least squares estimates. Standard errors are heteroskedasticity- and autocorrelation-consistent with a data-driven bandwidth equal to  $\text{floor}(1.75 \times T^{1/3})$ . Seven-day and thirty-day horizons are estimated on non-overlapping samples. The market

factor is the Bitcoin daily return. Volatility is the seven-day rolling standard deviation of Ethereum daily returns.

### Appendix A5. Neutral-handling policy and consolidated main-spec estimates

Policy and purpose. The main confirmatory specification excludes neutral headlines to reduce measurement noise in daily sentiment indices. This appendix consolidates the resulting estimates across assets and models. The mixed-source headline corpus (CoinDesk plus CryptoPanic API) cannot currently be re-queried in an identical way, and item-level provenance was not preserved during compilation. A numerical re-estimation with neutrals included is therefore not reproducible at present. Section 3.2.3 documents the construction and the intended robustness design for future replication when equivalent data access is available.

Estimation. OLS with Newey–West standard errors. Horizons are  $t = 0$  (same day),  $t = 7 D$ ,  $t = 30 D$ , and  $t + 1$ . The  $t$  statistic equals Coef divided by SE.  $\beta\_std$  is the standardized beta. Adj.  $R^2$  and  $N$  are taken from each regression.

#### Panel A. Consolidated main-spec results (neutrals excluded)

Model	Asset	Horizon	Coef	SE	t	p	$\beta\_std$	Adj. $R^2$	N
CryptoBERT	Bitcoin	$t = 0$	1.192	0.648	1.84	0.066	0.120	0.001	1335
CryptoBERT	Bitcoin	$t = 7 D$	0.251	2.545	0.10	0.921	0.025	-0.001	1335
CryptoBERT	Bitcoin	$t = 30 D$	0.651	6.412	0.10	0.919	0.065	-0.001	1335
CryptoBERT	Bitcoin	$t + 1$	-0.453	0.770	-0.59	0.556	0.045	0.001	1333
FinBERT	Bitcoin	$t = 0$	3.861	2.246	1.72	0.086	0.793	0.001	1297
FinBERT	Bitcoin	$t = 7 D$	1.785	5.311	0.34	0.737	0.367	-0.000	1297
FinBERT	Bitcoin	$t = 30 D$	2.885	3.019	0.96	0.339	0.593	-0.001	1297
FinBERT	Bitcoin	$t + 1$	2.360	1.947	1.21	0.226	0.485	0.005	1295

CryptoBERT	Ethereum	t = 0	0.670	0.801	0.84	0.403	0.085	-0.000	1213
CryptoBERT	Ethereum	t = 7 D	5.561	2.428	2.29	0.022	0.707	0.006	1213
CryptoBERT	Ethereum	t = 30 D	12.739	5.695	2.24	0.025	1.620	0.009	1213
CryptoBERT	Ethereum	t + 1	-0.211	0.780	- 0.27	0.786	- 0.028	-0.001	1211
FinBERT	Ethereum	t = 0	0.579	0.463	1.25	0.211	0.184	0.001	1202
FinBERT	Ethereum	t = 7 D	3.117	1.152	2.71	0.007	0.990	0.007	1196
FinBERT	Ethereum	t = 30 D	8.732	2.337	3.74	0.000	2.774	0.012	1173
FinBERT	Ethereum	t + 1	0.030	0.409	0.07	0.942	0.009	-0.001	1200

**Panel B. Neutrals included (within-day z scoring)**

Not available under current data access and provenance. The mixed-source headline file cannot be reconstructed with neutrals preserved at item level after CryptoPanic's API change and the loss of per-item source labels. The robustness recipe and all transformation steps needed for re-estimation are documented in Section 3.2.3 so that future researchers with equivalent access can reproduce this panel.