



Federal Reserve  
Bank of Dallas

# Social Network and Sentiment Contagion: Evidence from the Bitcoin Market

---

Bing Han, Haoyang Liu and Pengfei Sui

**Working Paper 2605**

**March 2026**

Research Department

<https://doi.org/10.24149/wp2605>

Working papers from the Federal Reserve Bank of Dallas are preliminary drafts circulated for professional comment. The views in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Dallas or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

# Social Network and Sentiment Contagion: Evidence from the Bitcoin Market\*

Bing Han<sup>†</sup>, Haoyang Liu<sup>‡</sup> and Pengfei Sui<sup>§</sup>

February 23, 2026

## Abstract

Using new data on social interactions and individual trading records in the Bitcoin market, we show that investor sentiment spreads across social connections. Investors systematically revise their beliefs about Bitcoin prices in the direction of average peer sentiment—even though that sentiment does not predict future prices. We document specific patterns in the diffusion of beliefs across networks, including evidence consistent with confirmation bias. Moreover, this social-sentiment contagion influences both individual trading decisions and overall market dynamics. Our novel measure of contagion intensity significantly forecasts Bitcoin volatility, trading volume, and market crashes.

**Keywords:** Social Interactions, Belief Updating, Sentiment Contagion, Bitcoin, Bubbles

**JEL Classification:** G11, G12, G41, G53

---

\* The authors thank Charles Cao, William Cong, Tony Cookson, Lawrence Jin, Weikai Li (discussant), Hongqi Liu, Yukun Liu, Song Ma, Marina Niessner, Jinfei Sheng (discussant), Tao Shu, Baolian Wang, Wei Xiong, Yao Zeng, Bohui Zhang, and the participants at the American Finance Association Annual Meetings, Annual Conference on Social and Behavioral Finance, China Financial Research Conference, GSU-RFS Fintech Conference, Five-Star Asia Pacific Workshop in Finance at Singapore Management University, Central University of Finance and Economics, Cheung Kong Graduate School of Business, Fudan University, Hong Kong University of Science and Technology, Nankai University, Nanjing University, National University of Singapore, Peking University Guanghua School of Management, Peking University HSBC Business School, Renmin University of China, Southern University of Science and Technology, University of Hong Kong, and University of International Business and Economics for their valuable comments on this paper. This work was previously titled “Conversations in the Market: Sentiment Contagion Among Investors” and “Social Learning and Sentiment Contagion in the Bitcoin Market”. The views expressed are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Dallas or the Federal Reserve System. All errors are our own.

<sup>†</sup>Corresponding author: Bing Han, University of Toronto, Bing.Han@rotman.utoronto.ca.

<sup>‡</sup>Haoyang Liu, The Federal Reserve Bank of Dallas, haoyang.liu@dal.frb.org.

<sup>§</sup>Pengfei Sui, The Chinese University of Hong Kong, Shenzhen, psui@cuhk.edu.cn.

# 1 Introduction

*“Man is by nature a social animal.”- Aristotle*

Investor sentiment has been a central theme throughout the history of modern finance. For decades, financial researchers have faced challenges in measuring investor sentiment and understanding its impact on markets. In recent years, significant progress has been made in studying investor sentiment at the aggregate market level. These studies have developed various sentiment measures to explain key market phenomena. Yet, an important question remains: how does investor sentiment evolve at the micro level? While extensive research has focused on aggregate investor sentiment, little attention has been given to the fundamental dynamics underlying these aggregate indicators.

In this paper, we fill this gap from a social perspective, by emphasizing the role of social interactions in belief formation. By investigating how individuals’ attitudes are affected by their peers’ discussions, we build a bottom-up approach to understand sentiment dynamics. We find that (i) social interactions lead to a spread of sentiment across investors – a pattern we define as *sentiment contagion*, (ii) the encountered peer sentiment in social interactions significantly predicts individuals’ future trading directions, and (iii) the intensity of sentiment contagion significantly predicts market outcomes, such as volume, volatility, and return crashes. Overall, our results highlight the pivotal role of social activities in understanding investor belief formation, trading decisions, and market dynamics.

To study the sentiment contagion via social interactions, we leverage data from an influential social platform in the Bitcoin market called *Bitcointalk*. Bitcointalk enables individuals to share opinions and engage with one another through posting and replying to others’ posts. One advantage of the Bitcointalk data is that it captures all conversations, providing a direct form of social interactions. This makes it an ideal environment for studying sentiment contagion, as it enables us to track individuals’ belief changes after engaging in conversations on social media. We determine the sentiment of each message using Natural Language Processing (NLP) techniques and examine changes in individuals’ post sentiment in response to peer feedback on the Bitcointalk forum. Furthermore, a subset of Bitcointalk users voluntarily disclose their Bitcoin wallet addresses to enhance the security of their forum accounts. Therefore, we can identify their trading records on the Bitcoin blockchain. By linking the sentiment expressed in their posts to their transactions, we can draw unique conclusions about the impact of sentiment contagion on investor trading behavior and market outcomes.

The Bitcoin market is also well-suited for our study due to its distinctive social characteristics. First, the development of Bitcoin is closely tied to social networks. Right after the inception of Bitcoin, its founder,

Satoshi Nakamoto, quickly established Bitcointalk, an online social platform designed to support Bitcoin operations and facilitate information sharing among Bitcoin users. Over time, Bitcointalk has grown into the most influential social platform on Bitcoin. Second, Bitcoin is a novel asset whose valuation is difficult to determine. People often turn to social platforms to learn about Bitcoin. As suggested by recent studies (Cong et al. (2021)), the valuation of Bitcoin depends on the aggregate demand which is a function of its social adoption. Thus, investors would find it useful to gather information from social media about others' views towards Bitcoin.

We begin by examining whether and how individuals' beliefs are affected by social interactions. The basic unit of observation in our analysis is a conversation segment that starts and ends with two consecutive posts (post[0] and post[1]) from the same individual. We use the term *social sentiment* to refer to the average sentiment of peers' messages between post[0] and post[1]. In our setting, sentiment contagion is manifested by one's belief update (measured by own sentiment change) after being exposed to peer messages. We find that an individual's sentiment moves in the same direction as the social sentiment, suggesting that peer views spread via conversations on social networks. Specifically, investors become more bullish (bearish) after being exposed to generally positive (negative) posts from others. Such sentiment contagion, or the updating of beliefs through social interactions, can lead to sentiment contagion.

We address several econometric issues when interpreting our results as evidence of sentiment contagion. First, we employ a test following Altonji et al. (2005) and Oster (2019) and demonstrate that the impact of endogeneity arising from omitted variables is limited. Our finding that social sentiment significantly affects belief updating remains robust in the presence of a wide range of controls. To account for common shocks that happen between the two posts, which may affect both the belief update of one poster and the sentiment of other posters, we control for the contemporaneous news sentiment from Ravenpack News Database, as well as the contemporaneous Bitcoin return and volatility. We also include lagged values of these control variables to account for possible delayed reactions to common shocks.

Second, we take into consideration sample selection related to posting on the social network. Our posting sample may disproportionately reflect the influence of frequent users, raising the question of whether our main finding merely reflects sentiment contagion among this specific group. However, we find that sentiment contagion remains strong regardless of the posting frequency. Another potential concern is the selective nature of posting a follow-up message after peer feedback. We demonstrate that, regardless of whether a follow-up message post[1] is present, the average peer sentiment over the next day after post[0] significantly predicts subsequent trading decisions of the original poster.

Our findings reveal that the contagion process exhibits pronounced heterogeneity across investor groups.

Less sophisticated investors, on average, respond more aggressively to social signals. For instance, after participating in a conversation with a social sentiment score of 1, inexperienced users exhibit an additional sentiment adjustment of 3.59%. Similarly, less informed users update their sentiment by an additional 2.66% compared to their informed counterparts under the same conditions. Furthermore, within a network structure – where a connection is defined by co-participation in at least one common conversation – we observe that more central individuals are significantly less responsive to social sentiment.

The contagion process is also characterized by distinct sentimental features. Peer sentiments expressed in post messages are largely uninformed and tend to negatively predict future Bitcoin returns. However, investors appear to perceive these social signals as informative, significantly updating their beliefs about Bitcoin prices in line with them. Notably, users also respond more strongly to social sentiment that aligns with their prior views, reflecting a clear consistency with confirmation bias. Sentiment contagion is unrelated to the emergence of Bitcoin-related news in traditional media. Instead, it intensifies on days with higher Bitcoin volatility or greater dispersion in RavenPack’s Bitcoin news sentiment, indicating a stronger reliance on social signals during periods of heightened uncertainty.

Taking a step further, we examine the impact of sentiment contagion on investors’ trading decisions. We establish a clear link between the social sentiment encountered by investors and their subsequent trading decisions for a subsample of Bitcointalk users whose trading records are available. Our findings demonstrate that peer sentiment from social interactions significantly predicts the direction of future trading. For one unit increase in the encountered social sentiment, investors are 0.41% more likely to buy Bitcoin the next day. The effect is economically significant, considering the unconditional probability of buying Bitcoin on a given day in our trading sample is 0.93%. Moreover, these transaction-based findings serve to validate our primary results about belief updating. At a broad level, these results establish a micro-foundation for understanding the influence of social interactions on market outcomes.

Furthermore, we document the impact of sentiment contagion on market dynamics. To measure the aggregate intensity of sentiment contagion, we count the number of Bitcointalk users who shift their sentiment in the same direction as peer sentiment during conversations, while controlling for time trends. As sentiment propagates through social interactions, it influences a growing number of investors, amplifies noise trading, and ultimately destabilizes the market. Therefore, the intensity of sentiment contagion should positively predict both Bitcoin trading volume and return volatility—a pattern we confirm in the data.

Finally, we demonstrate that sentiment contagion predicts future returns. Intuitively, investors may become either more optimistic or more pessimistic through sentiment contagion, exerting upward or downward noise trading pressure on Bitcoin prices, respectively. We hypothesize that the effect of sentiment contagion

on future returns depends on its cumulative net impact. For example, a positive cumulative net effect suggests that, on average, investors become more optimistic due to social interactions, exerting positive noise trading pressure and driving Bitcoin prices higher. However, this upward pressure is likely to result in subsequent price declines as the market undergoes a correction. Our findings strongly support this hypothesis: the cumulative net effect of sentiment contagion negatively predicts future returns and can even trigger extreme market movements, such as crashes. Overall, sentiment contagion is not merely a peripheral phenomenon; it plays a critical role in shaping both individual investment decisions and broader market dynamics.

Our study contributes to the literature on how investment ideas are transmitted, a foundational building block in the social finance literature. Early work by Shiller emphasizes the social nature of investment decision-making, with [Shiller et al. \(1984\)](#) describing investment as a social activity and [Shiller and Pound \(1989\)](#) identifying the transmission of financial information through social interactions. Building on this foundation, [Han et al. \(2022\)](#) develops a theoretical model to explore how investment ideas spread through social networks, while [Cookson et al. \(2024b\)](#) provides a comprehensive overview of the intersection between social media and finance. We present direct, granular evidence about how beliefs update and how sentiment spreads via social interactions.

Our paper also contributes to the growing literature on investor behavior in network settings, particularly by addressing the microfoundations of peer effects. Prior research has documented the influence of peer networks on investment decisions in various contexts ([Hong et al. \(2004\)](#); [Hong et al. \(2005\)](#); [Brown et al. \(2008\)](#); [Huang et al. \(2021\)](#)). For example, [Simon and Heimer \(2012\)](#) examine how interaction within investor networks promotes active trading, and [Heimer \(2016\)](#) find that social interactions amplify behavioral biases such as the disposition effect. [Chen et al. \(2014\)](#) demonstrate the usefulness of peer-based advice on SeekingAlpha. [Cookson et al. \(2023\)](#) document echo chamber effects on StockTwits, where selective exposure to reinforcing information drives trading volume and sustained disagreement. By documenting sentiment contagion resulting from social interactions and its impact on trading behavior, we advance the literature by identifying sentiment contagion as a key mechanism underlying peer effects and investor behavior in network settings.

One key question in the literature is whether peer effects wash out in aggregate. While empirical studies on this topic are relatively scarce, [Kuchler et al. \(2022\)](#) offer important evidence that social networks influence both investment decisions and firm-level outcomes, such as valuations and liquidity. [Cookson et al. \(2024a\)](#) further demonstrates that platform-level sentiment signals from investor social media (Twitter, StockTwits, and Seeking Alpha) predict stock returns and trading volume. Adopting a bottom-up approach, our study provides micro-level evidence connecting social interactions to aggregate market outcomes. We show that

sentiment contagion not only predicts individual trading decisions but also has significant implications for trading volume, volatility, and market crashes. This finding lays a critical foundation for understanding how individual social interactions aggregate to shape market dynamics.

Finally, we contribute to the growing literature on cryptocurrency markets, which has examined price discovery (Makarov and Schoar (2019)), arbitrage (Makarov and Schoar (2020)), and risk-return trade-offs (Liu et al. (2022)). We use the Bitcoin market as an ideal laboratory to study sentiment contagion along social networks, establishing significant connections between social interactions and market outcomes. Kogan et al. (2024) present a model where retail investors update their expectations based on recent price changes, which signal the likelihood of future adoption. In contrast, our study investigates the microfoundations of investor behavior, highlighting the role of peer sentiment – after controlling for past price changes – in shaping beliefs. By doing so, we establish a direct link between sentiment contagion and market dynamics, moving beyond price-based mechanisms.

The remainder of the paper is organized as follows: Section 2 describes the data and explains the main variables used in the analysis. Section 3 documents the sentiment contagion pattern. Section 4 studies the patterns of sentiment contagion. Section 5 presents the impact of sentiment contagion on individuals' trading decisions. Section 6 reports the impact of sentiment contagion on market outcomes. Section 7 concludes.

## 2 Data

### 2.1 Bitcointalk

Our study focuses on the Bitcoin market as it provides valuable opportunities to trace the full history of users' social interactive activities since its inception in 2009. To gather data on social interactions, we rely on Bitcointalk, which stands as a cornerstone of the cryptocurrency world, acting as the main online forum where Bitcoin investors, developers, and enthusiasts gather to share ideas and engage in discussion. The founder of Bitcointalk, Satoshi Nakamoto, is the presumed pseudonymous author (or authors) of the original Bitcoin white paper, which describes Bitcoin's reference implementation. Its popularity has been widely covered by influential media outlets such as the Wall Street Journal, Forbes, and Bloomberg. Considering its impact on Bitcoin investors, conversations on this forum serve as a good proxy for social interactions in the Bitcoin market.

In this paper, our research is centered around posts on Bitcointalk, which are timestamped messages containing sentences written by users. Generally, sentences within each post are concise in nature. These

posts are organized within threads, which are collections of sequential posts that revolve around a common theme, allowing users to interact with each other. After posting in one thread, the user receives notifications from the forum if others publish follow-up posts in the same thread. Therefore, users are likely to pay attention to subsequent posts in the thread. As an example in our sample, Figure 1 presents a thread titled “What does the future of bitcoin look like” joined by user DavidLuziz on May 31, 2018, with subsequent participation from other users sharing their opinions.

On Bitcointalk, threads are displayed under different “child boards”, which are subforums that gather discussions related to specific topics. Not all subforums are about topics related to the prospect of Bitcoin prices. For instance, the “Bitcoin Technical Support” child board is dedicated to addressing technical questions concerning Bitcoin Core, nodes, the Bitcoin network, transactions, and addresses. For our study, we specifically focus on the “Speculation” child board, where users predominantly discuss their views on Bitcoin’s prospects. Figure 2 presents the visual representation of the keywords mentioned in posts within this child board. The most frequent words include “Bitcoin”, “price”, “think”, and “know”. These keywords indicate that the topics are quite concentrated, and investors express their attitudes about Bitcoin prices in this subforum. The “Speculation” child board serves as an ideal field laboratory for our research, as we can examine how social interactions propagate users’ views about Bitcoin, and how sentiment contagion affects individual trading behavior and aggregate market outcomes.

There are multiple layers of user heterogeneity on this forum. For instance, we observe variations in the level of sophistication among users. Bitcointalk utilizes a merit system wherein users are awarded points for their contributions to the Bitcoin community. Based on their merits and activities, users are assigned different ranks on the Bitcointalk website, ranging from the high rank “Legendary” to the low rank “Newbie”. Typically, users with higher ranks demonstrate a greater understanding of Bitcoin.

## 2.2 Textual Analysis and Post Sentiment

We measure Bitcointalk users’ beliefs about future Bitcoin prices by analyzing the sentiment of their posts. We develop a methodology to extract user sentiment from posts on the Bitcointalk forum. Since our raw data consist of over 1 million posts – far beyond what we can manually interpret – we employ a textual analysis algorithm. Our algorithm is a two-step procedure based on a keyword dictionary and a natural language processing (NLP) algorithm developed by the Stanford University NLP group (Manning et al. (2014)).<sup>1</sup>

---

<sup>1</sup>While methods based on large language models (LLMs) are capable of extracting rich information from text, they may encounter potential look-ahead biases in certain contexts. To address this, our approach is specifically designed to mitigate such risks. As

In the first step, we randomly select 10,000 sentences from our dataset and manually label them into four categories: “positive”, “neutral”, “negative”, and “irrelevant”. This approach mainly follows the methods used in Baker et al. (2016) and Tetlock (2007). We then use these labeled sentences as our training set and construct a keyword dictionary for each category. For example, keywords under the “positive” category include “buy”, “increase”, and “rise”, while the “negative” category contains phrases such as “sell”, “decrease”, and “plunge”, and so on. Additionally, we process sentences containing negative particles like “not” and “couldn’t”. If a sentence with negative particles contains words that are classified as “positive”, we reverse our interpretation and label the sentence as “negative” instead of “positive”. The dictionary for the “neutral” category includes terms such as “hold”, “wait”, “unpredictable”, and others. If a sentence does not contain any of the keywords in the above three categories, we label it as “irrelevant”.

In the second step, we apply the Stanford NLP algorithm to detect sentences that describe the future. One challenge we face in our labeling exercise is distinguishing between descriptive statements about past performance and beliefs about the future in certain posts. For instance, a sentence like “Bitcoin market really increased a lot in the past months” can be ambiguous in terms of an user’s opinion about future Bitcoin prices. Relying on the Stanford NLP algorithm, we can identify sentences with forward-looking statements, and label the rest as “irrelevant”. For example, if a sentence contains the keyword “increased”, the Stanford NLP would detect the past tense, and our algorithm would label it as “irrelevant”. On the other hand, the Stanford NLP algorithm would identify future tense in sentences with phrases such as “will increase”. Furthermore, our definition of forward-looking statements aligns with the format recommended by the SEC for 10-K reports, including sentences with terms like “expects”, “anticipates”, and so on.<sup>2</sup>

After the aforementioned two steps, we extract user sentiment from the relevant sentences. We assign a value of 1 to sentences with a “positive” sentiment, -1 for “negative”, and 0 for all others. For example, a sentence like “Bitcoin price will roar” is assigned a value of 1, while “Bitcoin price is not going to fall” is also assigned a value of 1. Conversely, a sentence like “Bitcoin is doomed to fall” is assigned a value of -1, and a sentence like “Bitcoin price is really unpredictable” takes a value of 0. To calculate the overall sentiment of a post, we take the average of the sentiment values of all its sentences.<sup>3</sup> As a result, the sentiment measurement for a post falls within a continuous range from -1 to 1. The out-of-sample accuracy of our algorithm is approximately 85%.

---

a robustness test, we employ LLMs to assign continuous sentiment scores at the post level, confirming that our findings remain robust. The details of this LLM-based method are provided in Section A. Table IA.1 summarizes our primary findings using these LLM-assigned sentiment scores.

<sup>2</sup>Please see [https://www.sec.gov/Archives/edgar/data/1082027/000139390519000101/neik\\_10k.htm](https://www.sec.gov/Archives/edgar/data/1082027/000139390519000101/neik_10k.htm) for details.

<sup>3</sup>In Appendix Tables IA.3 and IA.4, we use the maximum and the sum of the sentiment values from all sentences. Our findings remain robust.

## 2.3 Transaction Data

We are able to link the sentiment of a subsample of users to their trading behaviors. This group of users voluntarily published their Bitcoin wallet addresses in a thread officially initiated by the forum organizer, as a protective measure for their Bitcointalk accounts against hacking.<sup>4</sup> Since Bitcoin transactions are publicly available on the Bitcoin Blockchain, we can trace their detailed transaction information, such as trade size and trade timestamps, using these published wallet addresses.

## 2.4 Other Data Sources

To proxy investor attention for the Bitcoin market, we utilize Google search volume data following [Dat et al. \(2011\)](#). Specifically, we define Google Search (Bitcoin) as the difference between the Google Search Volume Index and its past one-month mean, divided by the lagged one-month mean.

To systematically capture the impact of news events on the Bitcoin market, we obtain news related to Bitcoin from Ravenpack News Analytics, a reliable source that tracks news reports about Bitcoin and provides sentiment scores for them dating back to 2011. We refer to the general tone of news coverage on Bitcoin in the media as the news sentiment, and include it as a control variable in our analysis.

We collect Bitcoin market data, including hourly returns and trading volume, from CoinAPI. We also calculate daily return volatility. This data is also utilized in [Griffin and Shams \(2020\)](#). The transactions before May 2012 were relatively less frequent. Therefore, our focus is on transaction data starting from May 1st, 2012, and the sample period ends on July 30th, 2022.

## 2.5 Summary Statistics

To align with the Bitcoin market data, our main Bitcointalk dataset covers the period between May 1, 2012, and July 30, 2022. It comprises 666,006 posts in the "Speculation" subforum contributed by 44,356 users. Among these users, we successfully linked 2,550 individuals to their transaction records via their disclosed Bitcoin wallet addresses. In Panels A, B, and C, we present summary statistics for user sentiment at the user, daily, and thread levels, respectively. Panel D reports summary statistics for Bitcoin market information, RavenPack news sentiment, and Google Search Volume.

---

<sup>4</sup>This is an official activity organized by the management team of the Bitcointalk forum. Users of Bitcointalk are recommended to disclose their public Bitcoin wallet address within this thread. When a user posts their public Bitcoin wallet address on Bitcointalk, they are creating a public record of their ownership of that wallet address. If their Bitcointalk account is later hacked, they can use their private key to create a digital signature that proves their ownership of the Bitcoin address they previously shared. This digital signature can then be used to authenticate the account holder's identity and facilitate the recovery of their Bitcointalk account.

In Panel A of Table 1, we present the summary statistics for posting activities at the user level. For a representative user in our sample, the average sentiment score of the posts they published is 0.32, suggesting a general optimism towards Bitcoin among users. Such optimism is consistent with the existing literature on motivated beliefs (Bénabou and Tirole (2016)) and utility-based biases in beliefs (Brunnermeier and Parker (2005)). Furthermore, we observe a significant within-user standard deviation of sentiment at 0.61, indicating substantial variation in sentiment over time at the individual level. We require at least two posts for a given user to compute the standard deviation of post sentiment. The distribution of the number of posts is skewed, as evident from the difference between the median and mean number of posts per user.

Panel B of Table 1 describes post activities on a given day. On a representative trading day, about 120 users publish around 178 posts. The average post sentiment is 0.28, which is optimistic and consistent with the statistics at the user level.

Panel C reveals that posts within threads often exhibit dispersed sentiment, with an average within-thread standard deviation of 0.65. Furthermore, active user engagement is notable, with 25 users participating in a representative thread containing a total of 35 published posts. Overall, users seem to pay close attention to the posts, resulting in active exchanges within threads, with a median gap between consecutive posts within a thread of only 48 minutes.

Panel D provides summary statistics for variables at the aggregate market level. In the first row, the mean return (annualized) is 119.0% and the standard deviation is 16.17. When combined, the Sharpe ratio is approximately 0.074. The second row presents the within-day volatility of Bitcoin, calculated using hourly return data. On average, the volatility is approximately 3.7%. We present volume-related variables for Bitcoin, spanning from the third to the fifth row. On a typical trading day, there are 16,472 transactions involving 9,359 bitcoins being traded. The average daily trading volume totals 56.5 million dollars. Moving to the sixth row, we present summary statistics for the sentiment of RavenPack News. In its original form, the sentiment score for RavenPack News Database ranges from 0 to 100, but we have normalized it to the range of  $[-1, 1]$ . On average, the transformed sentiment of news within a single day is 0.041 with a standard deviation of 0.330.

### 3 Evidence for Sentiment Contagion

In this section, we investigate whether peer sentiment spread via social interactions, a pattern we refer to as *sentiment contagion*. Specifically, we examine how a user's belief may be influenced by the opinions of others, as expressed through conversations on the forum.

### 3.1 Data Structure

The unit of analysis in our study of sentiment contagion is what we term a *conversation*: an episode of social interactions on the Bitcointalk forum. It consists of a pair of consecutive posts by a user, along with comments from other users in between. Figure 1 illustrates an example of a conversation. User DavidLuziz published one post at 05:37:35 AM on May 31st, 2018. Subsequently, several other users joined the conversation and shared their views about Bitcoin. At 07:50:10 AM on the same day, DavidLuziz published another post in the same thread, becoming more optimistic about Bitcoin after interacting with other users.

Figure 3 illustrates the timeline of our analysis. For a pair of consecutive posts by a given user, we denote the first post published at timestamp  $t_0$  as `post[0]`, and the second post made at timestamp  $t_1$  as `post[1]`.<sup>5</sup> There is a trade-off when determining the time window  $[t_0, t_1]$  of a conversation. A shorter time window would reduce the confounding influence of unobservable factors for belief changes, but it would result in a smaller sample of conversations, thereby reducing the statistical power of our tests. We strike a balance by choosing a window that is sufficiently short in order to cleanly capture belief changes caused by social interactions while avoiding excessive data loss. Figure 4 presents the cumulative distribution function of the time gap between consecutive posts by the same user in our sample. More than 50% of the pairs fall within a 24-hour window. Therefore, we select a 24-hour time window in our main analysis.<sup>6</sup> In total, there are 296,291 pairs of posts by the same user (and, correspondingly 296,291 conversations) from 16,535 users in our main test sample.

### 3.2 Empirical Strategy

To test sentiment contagion via social interactions on Bitcointalk, we examine whether and how a user updates her belief about Bitcoin after being exposed to others' views in the conversation. We refer to the sentiment in `post[0]` and `post[1]` as the *prior sentiment* and the *posterior sentiment*, respectively. We link the user's sentiment change – defined as the revision from her prior to posterior sentiment – to the *social sentiment*, defined as the average sentiment of other users' posts within the conversation. We control for confounding factors that might affect sentiment change, such as Bitcoin return and volatility, common news,

---

<sup>5</sup>One user's consecutive posts could appear in two different threads. In such cases, conversations refer to posts published between  $t_0$  and  $t_1$  in both threads to which `post[0]` and `post[1]` belong. Our results remain robust when restricting the sample to cases where `post[0]` and `post[1]` are in the same thread (see Appendix Table IA.5).

<sup>6</sup>Appendix Table IA.6 indicates that the effect of social sentiment on belief updating is even stronger for 48- or 72-hour windows than for the 24-hour window.

and the overall sentiment of contemporaneous posts on the forum. Our main regression model is:

$$\begin{aligned} \text{Senti Change}_{i,j,t_0 \rightarrow t_1} &= \beta_1 \text{Social Sentiment}_{i,j,t_0 \rightarrow t_1} \\ &+ \gamma \text{Control}_{i,t_0 \rightarrow t_1} + \text{Fixed Effects} + u_{i,t_1}, \end{aligned} \quad (1)$$

where  $\text{Senti Change}_{i,j,t_0 \rightarrow t_1}$  is the revision in user  $i$ 's sentiment from post[0] at  $t_0$  to post[1] at  $t_1$  during conversation  $j$ .<sup>7</sup>  $\text{Social Sentiment}_{i,j,t_0 \rightarrow t_1}$  is the average peer sentiment that user  $i$  has encountered in conversation  $j$  between time  $t_0$  and  $t_1$ . We include a set of control variables, denoted as  $\text{Control}_{i,t_0 \rightarrow t_1}$ , along with various fixed effects to address several concerns related to issues such as endogeneity. We report standard errors that are clustered at the user and day levels.

Sentiment change, the dependent variable in the regression 1, captures the effect of sentiment contagion and aligns with our subsequent analysis of the impact of sentiment contagion on individuals' trading and market outcomes. An alternative dependent variable is the ex-post sentiment. Appendix Table IA.2 verifies the robustness of our findings under this alternative setting. Moreover, our regression models incorporate fixed effects to account for the signs of the prior sentiment (hereafter *prior fixed effect*). This choice is motivated by the observation that users with a positive prior sentiment naturally have less room for increased optimism compared to those with a negative prior sentiment.

### 3.3 Evidence for Sentiment Contagion

The baseline regression in Table 2 column (1) shows a significantly positive relation between sentiment change and social sentiment with a t-statistic of 6.81 for the regression coefficient of social sentiment. A one standard deviation increase in social sentiment (42.18%) is associated with an increase in changes in sentiment of 1.20% (calculated as  $42.18\% \times 2.84\%$ ).

Columns (2) to (5) show that the effect of social sentiment remains significant and robust after we include various controls to address potential concerns about omitted variables. Column (2) adds the average news sentiment from Ravenpack News Database between consecutive posts to control for contemporaneous common news arrivals (Feng and Seasholes (2004)). In column (3), we include the Bitcoin return and volatility between  $t_0$  and  $t_1$ . In column (4), we further control for the contemporaneous forum sentiment, computed as the average sentiment of posts published between  $t_0$  and  $t_1$  in all conversations on Bitcointalk other than those involving post[0] or post[1]. Furthermore, to address concerns that investors may respond to news with a delay, we include in column (5) lagged Ravenpack news shocks up to 7 days before post[0] as additional controls. In Appendix Table IA.7, we further include Bitcoin returns and volatility from the past

<sup>7</sup>Appendix Table IA.2 verifies the robustness of our findings when we use the posterior sentiment as the dependent variable.

24-hour, 48-hour, and 7-day windows, as well as the average sentiment of posts before post[0], to account for investors' delayed responses to market dynamics and social interactions. We find that the regression coefficient of social sentiment remains significantly positive in all cases. We also verify the robustness of our results when user-week fixed effects are included (so that we solely rely on variation between different post pairs by the same user in the same week), suggesting that our results are not driven by time-varying user characteristics.

### 3.4 Placebo Test

To further establish the significance of social interactions for belief updating, we conduct a placebo test. Specifically, we investigate whether a Bitcointalk user's belief is influenced by the average sentiment of randomly selected conversations that take place between the timestamps of post[0] and post[1] in which the user does not participate. The results are presented in Table 3. Column (1) shows that an individual's sentiment about Bitcoin price is not significantly affected by the average sentiment of other users in a randomly selected conversation in which the individual did not participate. This pattern remains robust when we select multiple random conversations. For example, in column (4), we randomly select ten conversations that occur between post[0] and post[1], in which the user does not participate. We continue to find that an individual's belief is not affected by the views of others to whom they are not linked. The result of this placebo test underscores the importance of social interactions in influencing belief changes.

### 3.5 Endogeneity

We address potential endogeneity concerns in our tests. Reverse causality is unlikely in our setting, as we examine whether one's belief changes significantly *after* being exposed to others' views in a conversation. Another concern is "omitted variable bias" (Jiang (2017)). It is possible that the belief updating of a user and social sentiment are both affected by some unobserved factor, such as the arrival of public news. One advantage of our analysis is that we observe the timestamp of users' posts. This enables us to construct granular controls such as concurrent news arrivals and market returns. For example, we include the average news sentiment from the Ravenpack News Database between the two consecutive posts as a regressor to control for the effect of contemporaneous news arrivals (Feng and Seasholes (2004)). We further account for the Bitcoin returns between the two consecutive posts in order to absorb the possible effect of common shocks that occur between post[0] and post[1]. Another source of omitted variables arises from users' lagged responses to events that occur before the conversation starts. Since we observe the exact timestamp when a

conversation starts, we can address this issue by including lagged news arrivals, lagged market returns, and lagged posting activities.

Despite our inclusion of nuanced controls to mitigate the omitted variable issues as much as possible, concerns may remain that these controls do not fully address endogeneity. To delve deeper and evaluate the severity of omitted variable bias, we conduct a test following [Oster \(2019\)](#). The idea behind the test is that it is important to interpret movements in coefficients and  $R^2$  jointly. If the existing controls significantly improve  $R^2$  without moving the coefficient much towards 0, then the estimated effect is probably not due to an omitted variable. We can extrapolate and infer that adding more controls will not change the results to be statistically insignificant.

The formal test incorporates the change in  $R^2$  resulting from the addition of controls. To conduct this test, we first estimate the linear model presented in equation (1) without any control variables, allowing us to obtain the coefficient of social sentiment  $\beta_u$  and the coefficient of determination  $R_u^2$ . Then we include contemporaneous and lagged control variables (see the last column of Table 2) to derive the coefficient  $\beta_c$  and the coefficient of determination  $R_c^2$ .

[Oster \(2019\)](#) shows that under some technical conditions, the null hypothesis of no relation can be tested by checking whether zero belongs to an identified set of coefficients defined as the interval between  $\beta_{adj}$  and  $\beta_c$ , where the bias-adjusted coefficient  $\beta_{adj}$  is given by

$$\beta_{adj} \approx \beta_c - \delta \frac{(\beta_u - \beta_c)(R_{max}^2 - R_c^2)}{R_c^2 - R_u^2}.$$

The identified set of coefficients relies on two crucial parameters:  $\delta$  and  $R_{max}^2$ . The parameter  $\delta$  quantifies the level of selection on unobservable factors relative to observable controls. A higher value of  $\delta$  indicates a greater severity of the omitted variable problem. On the other hand, the parameter  $R_{max}^2$  represents the hypothetical overall  $R^2$  of the model when both observable and unobservable variables are accounted for. It indicates the extent to which the variation in the outcome variable can be explained by controlling for all relevant factors. An upper bound on  $R_{max}^2$  is 1. [Oster \(2019\)](#) argues that a reasonable case to consider is  $\delta = 1$  and  $R_{max}^2 = 1$ . In our analysis, we experiment with different parameter values.

Table 4 presents the results of the analysis. In the upper left panel, we follow [Mian and Sufi \(2014\)](#) and use  $R_{max}^2 = \min(2.2R_c^2, 1)$  and  $\delta = 1$ . The adjusted  $\beta$  value is 0.083, allowing us to reject the null hypothesis that  $\beta = 0$ . Moving to the right panel, we adopt a more aggressive stance with  $\delta = 2$ , corresponding to the unrestricted estimator in [Oster \(2019\)](#). Despite this, the identified set is  $[0.031, 0.177]$  which remains significantly distant from zero. In the bottom panel, we set  $R_{max}^2$  to be 1, and regardless of whether  $\delta$  equals 1 or 2, we can still reject the null hypothesis. Overall, results in Table 4 show that across a wide range

of scenarios, omitted variables would not qualitatively affect the significance of the positive relationship between social sentiment and sentiment change.

Since participating in a conversation is voluntary, another concern is that individuals may be more inclined to sustain their involvement in conversations that align with their prior sentiment. When the follow-up messages mostly differ in sentiment from an individual’s initial post (post[0]), the user might choose to leave the conversation (i.e., there is no post[1]).<sup>8</sup> Such scenarios are not included in our tests (because we require post[1] to exist in order to measure a user’s belief update). If there is no sentiment contagion in such scenarios (e.g., people who refrain from posting messages to debate with different views are stubborn or confident about their own prior, so they do not change their views in the presence of social signals), then our test would produce an overestimation of the intensity of sentiment contagion for the general population. However, our evidence of the influence of peer sentiment on belief updating remains valid for users with both post[0] and post[1].

We provide additional evidence in support of sentiment contagion in a subsample of users without follow-up messages. Specifically, we leverage the users’ trading data to measure changes in belief. For users who voluntarily reveal their wallet address<sup>9</sup>, their trading records are observable. For posts by these users, we measure social sentiment as the average sentiment of posts within a subsequent 24-hour window after post[0], but we do not require the presence of a follow-up post[1]. Instead of linking social sentiment to subsequent sentiment change, we examine whether social sentiment positively predicts future trading decisions. This is confirmed in the data (see Section 5 and Table 11 for details), supporting the idea that social signals significantly impact users’ belief updating and trading.

We also examine whether our main finding merely reflects sentiment contagion among frequent posters as our sample observations come disproportionately from users who choose to post messages frequently. We categorize users based on their posting frequencies into two subsamples and rerun our main regressions separately. Specifically, we calculate the average time interval between users’ consecutive pairs of posts. We define users with an average interval above the population average as infrequent posters, and the remaining users as frequent posters. Our findings are presented in Table 5. Columns (1) and (2) indicate a strong and significant sentiment contagion for both the frequent and infrequent users. Furthermore, Column (3) shows that there is no significant difference in the intensity of sentiment contagion between frequent and infrequent posters.

---

<sup>8</sup>A plausible explanation for such behavior is “emotion regulation” (see Hwang (2023)).

<sup>9</sup>For detailed discussions of this subgroup of users, please see Sections 2.4 and 5.

## 4 Evidence of Non-fully Rational Belief Updating in Social Interactions

In this section, we investigate the nature of peer signals and belief updating during social interactions. Recent studies reveal systematic deviations from rational Bayesian updating to signals, such as the over-weighting of certain information, selective exposure, and reactions to uninformative cues. We show that peer sentiment expressed in posts is largely uninformed—it negatively predicts future Bitcoin returns—yet investors treat it as informative and respond to it significantly. Moreover, we document heterogeneity in the strength of our main findings across user groups and over different time periods, which is consistent with the existence of non-fully rational components in social sentiment contagion. Finally, our analysis highlights the role of specific cognitive biases in belief updating via social interactions, such as the availability heuristic and confirmation bias.

### 4.1 Informativeness of Social Sentiment

We start by investigating the informativeness of peer sentiment on the Bitcointalk forum. We estimate the relationship between peer sentiment and future Bitcoin returns by running the following regression:

$$\text{Cumulative Return}_{i,j,t_1+1 \rightarrow t_1+k} = \beta_1 \text{Social Sentiment}_{i,j,t_0 \rightarrow t_1} + \gamma \text{Control}_{t_0 \rightarrow t_1} + \text{Fixed Effects} + u_{i,t_1},$$

where the dependent variable, denoted as  $\text{Cumulative Return}_{i,j,t_1+1 \rightarrow t_1+k}$ , represents the cumulative return in a forward-looking window from hour 1 to hour  $k$  after user  $i$  participates in a conversation  $j$  that ends at timestamp  $t_1$ . We study five different time windows: 6, 24, 48, and 72 hours, as well as seven days after the conversation. The main explanatory variable,  $\text{Social Sentiment}_{i,j,t_0 \rightarrow t_1}$ , represents the average peer sentiment encountered by user  $i$  in conversation  $j$  between timestamps  $t_0$  and  $t_1$ . To address potential confounding factors, we further include a set of time-varying controls. We also include user and date fixed effects, along with the prior fixed effect. Standard errors are clustered by user and day, as recommended by [Hodrick \(1992\)](#).

This specification follows a well-established approach in the retail trading literature ([Barber and Odean \(2000\)](#)), which evaluates the quality of traders' decisions by comparing the timing of their buys and sells to subsequent returns. Similarly, we evaluate the quality of social signal by testing the predictive power of social sentiment for subsequent Bitcoin returns. A positive coefficient for  $\beta_1$  in equation (??) would indicate that social sentiment is informative and thus worth incorporating. However, [Table 6](#) shows that social sentiment is largely misinformed as it negatively predicts future returns over most horizons. For example, a strategy that buys Bitcoin immediately after observing a social sentiment of 1 and holds for the next 24 hours would, on average, incur a loss of -0.094%. This suggests that users treat peer opinions encountered on social media

as if they were useful and informative. They put a positive weight on the social signal and update their beliefs towards the peer sentiment, even though, on average, it does not correctly predict future Bitcoin price movements.

While, on average, peer sentiment does not positively forecast returns, there may still be some users whose posts contain useful signals about future Bitcoin direction. Users may listen more to such posters. To explore this hypothesis, on a rolling basis, we define less informed users as those whose sentiment exhibits a low correlation with future Bitcoin returns over the next 7 days (correlation < 25th percentile). We also look for other types of posters that are potentially more sophisticated or experienced with the Bitcoin market. For example, users on Bitcointalk are classified into legendary and newbie categories based on their influence and contribution to the Bitcointalk community. We designate the newbie category as naïve users and legendary category as sophisticated users. Another example is central users (“influencers”) who post widely across many conversations. We treat two users as connected if they have participated in at least one conversation together. Following an established algorithm in the network literature (Hagberg et al. (2008)), we calculate the centrality score for each user on a daily basis. For each category of users, we investigate separately whether each group’s sentiment expressed in social interactions predicts future Bitcoin returns and whether users listen to their views.

Panel A of Table 7 presents our findings. In columns (1) and (2), we compare the informativeness of social sentiment from legendary and naive users. Although social sentiment from naive users is slightly less informative than that from legendary users, both negatively predict cumulative returns over the next 24 hours. Columns (3) and (4) demonstrate that the social sentiment of both central and non-central users negatively predicts returns. Somewhat surprisingly, social sentiment from central users more strongly predicts negative returns than social sentiment from non-central users. Columns (5) and (6) show that social sentiment from both informed and less informed users negatively predicts returns, although social sentiment from informed users seems relatively more precise.

In Panel B of Table 7, we investigate how users respond to social sentiment from users with different features. In columns (1) and (2), we find that users are twice as sensitive to social sentiment from naive users as to that from legendary users. This finding is difficult to reconcile with a rational narrative, as social sentiment from both naive and legendary users negatively predicts future returns. As shown in columns (3) and (4), users strongly respond to social sentiment from both central and non-central users. However, social sentiment from both types of users negatively predicts future returns, as shown in columns (3) and (4) of Panel A. Columns (5) and (6) indicate that users respond positively to social sentiment from less informed users and informed users, but this social sentiment also negatively predicts future returns. Overall,

the findings from Table 6 and Table 7 provide support for a behavioral narrative of sentiment contagion: peer sentiment is generally uninformative, but users significantly revise their beliefs in response to peer opinions from social interactions.

## 4.2 Heterogeneity across Users

Next, we investigate which types of users react more strongly to social signals. First, we examine how sentiment contagion differs between naïve users and sophisticated users. Column (1) of Table 8 presents the heterogeneity in sentiment contagion across these two groups. The interaction term between the naïve user indicator and social sentiment is positive and significant, suggesting that naïve users are more responsive to social sentiment compared to sophisticated users. On average, when faced with a social sentiment of 1, a naïve user changes own sentiment by 3.585% more than a sophisticated user. Similarly, Column (3) of Table 8 indicates that, compared to more informed users, less informed users update their beliefs by 2.66% more after participating in a conversation with a social sentiment of 1.

We also compare the effects of peer sentiment on central posters in the Bitcointalk network versus other users. Column (2) of Table 8 shows that central users respond significantly less to peer sentiment. A one-unit increase in the centrality score reduces the magnitude of sentiment contagion by 0.047%. The sentiment contagion effect is moderated for central users, but it does not get eliminated. For instance, even for a user at the 75th percentile with a centrality score of 16.3, the effect of sentiment contagion remains at 3.32% (calculated as  $4.09\% - 0.047\% \times 16.3$ ).

Overall, there is substantial heterogeneity in the strength of social sentiment contagion. Less sophisticated, less experienced, and less informed investors—those who are less rational—are more responsive to peer sentiment. These findings are consistent with a non-fully rational mechanism underlying sentiment contagion in social interactions.

## 4.3 Time-Variation in Social Sentiment Contagion

Table 9 examines how the strength of social sentiment contagion varies over time. First, we investigate whether sentiment contagion is related to information flow. We identify a day as informative if the total number of Bitcoin news arrivals in the RavenPack News database exceeds its sample median.<sup>10</sup> We expect stronger belief updating on informative days if posts on Bitcointalk mainly serve to propagate the public news (e.g., by calling attention to the news). Column (1) of Table 9 indicates a statistically insignificant

---

<sup>10</sup>In untabulated analysis, we verify that our findings remain the same when we classify a day as informative if the total number of “novel” news items labeled by RavenPack exceeds its sample median.

coefficient for the interaction term of social sentiment and the dummy for informative days, suggesting that sentiment contagion is not stronger when there is more news. Therefore, the social sentiment contagion we document does not seem to be driven by news flow.

Second, we investigate whether social sentiment contagion is influenced by the level of uncertainty about Bitcoin. [Tversky and Kahneman \(1974\)](#) established that the availability heuristic influences judgments made under uncertainty. In our context, during periods of heightened uncertainty, the availability heuristic may prompt individuals to be more inclined to conform to the social sentiment they encounter. Column (2) of Table 9 shows that sentiment contagion is significantly stronger when the standard deviation of news sentiment on Bitcoin from the RavenPack News database is higher than the sample median. We observe the same pattern in column (3), where we proxy for uncertainty using the standard deviation of hourly Bitcoin returns in a rolling 24-hour window.

Overall, we provide consistent evidence for elevated sentiment contagion in uncertain periods. In contrast, rational Bayesian learning should be weaker during periods of higher uncertainty when the signals are noisier. One behavioral explanation for our finding is the availability heuristic, which influences judgments made under uncertainty ([Tversky and Kahneman \(1974\)](#)): during periods of heightened uncertainty, the availability heuristic may prompt individuals to rely more on the peer sentiment they encounter.

#### 4.4 Asymmetry and Confirmation Bias

The psychology literature has extensively documented humans' cognitive limitations, which affect our perception, attention, memory, decision-making, and problem-solving abilities. One notable heuristic is confirmation bias ([Nickerson \(1998\)](#)), the tendency to search for, interpret, favor, and recall information in a way that confirms or supports one's prior beliefs or values.

Confirmation bias can lead users to interpret information in a way that confirms their priors. To examine this phenomenon, we investigate how users selectively interpret signals they encounter in conversations. Specifically, we create two new variables: social sentiment (-) and social sentiment(+). Social sentiment (-) is defined as the minimum value of social sentiment and zero. Social sentiment (+) is defined as the maximum value of social sentiment and zero. These two variables capture the linear relationship between sentiment contagion and social sentiment in the negative and positive regions, respectively. We then conduct a regression test of belief updating similar to Equation 1 using two separate subsamples: one for users with positive priors and one for those with non-positive priors. Confirmation bias predicts that users with positive priors would respond more strongly to social sentiment (+), while users with negative priors would respond more aggressively to social sentiment (-).

Table 10 presents our findings. In column (1), we observe that users with positive priors adjust their sentiment towards the social sentiment, which supports the main message conveyed in Table 2. Column (2) shows that users with positive priors respond to both positive and negative social sentiment, but they are much more responsive to positive sentiment. Specifically, a one-unit increase in positive social sentiment is associated with a 5.012% probability (statistically significant at the 1% level) of a sentiment upward change. In contrast, a one-unit increase in negative social sentiment is associated with only 2.687% probability (which is not significant at the 10% level) of own sentiment change. In column (3), we find that users with negative priors also adjust their sentiment towards the social sentiment. Finally, in column (4), we observe that users with negative priors respond to both positive and negative social sentiment, but they are more responsive to negative social sentiment. In untabulated tables, we include the user-week interactive fixed effect and find that our results remain similar. Taken together, the results in Table 10 indicate an asymmetric pattern in the belief updating via social interactions that is consistent with confirmation bias: investors update their beliefs in the direction of social sentiment more aggressively when the latter is consistent with their priors.

## 5 Sentiment Contagion and Individual Trading

In this section, we provide micro-level evidence on the link between sentiment contagion and individuals' trading decisions for a subsample of Bitcointalk users. Out of the 44,356 registered users in our sample, 2,550 users voluntarily published their Bitcoin wallet addresses, following the recommendation of the Bitcointalk website (see footnote 3). Since Bitcoin transactions are public on the blockchain, we can identify the complete transaction history of these 2,550 users. Notably, regardless of whether users choose to reveal their wallet address, the main effect we document in Table 2 remains robust. The intensity of sentiment contagion exhibits no significant difference between the two subsamples, as shown in column (3) of the Appendix Table IA.8. Furthermore, the transactions of these 2,550 Bitcointalk users are representative of the general trading pattern in the Bitcoin market. The total trading volume by these investors closely tracks the total dollar volume of Bitcoin, with a correlation coefficient of 0.46.

To establish the link between sentiment contagion and individual trading, we use the peer sentiment from the conversation on Bitcointalk that an individual is involved in to predict her Bitcoin transactions the next day. Specifically, we run the following regression:

$$\text{Net Buy Dummy}_{i,t+1} = \beta \text{I(Social Sentiment)}_{i,t} + \gamma \text{Control}_{i,t} + \text{Fixed Effects} + u_{t+1}. \quad (2)$$

Given the relatively small number of users with observed Bitcoin transactions, we no longer require the users

to have consecutive posts within a 24-hour window in the above test in order to increase the statistical power. On each day, we focus on users who have made at least one post over the past 7 days. Otherwise, the user-day pairs are excluded from the above regression. We assume that when individuals participated on Bitcointalk, they were exposed to other users' subsequent posts within a 24-hour window in the same thread. We analyze all the threads in which an individual has participated over the past 7 days and define social sentiment as the average sentiment of other users' posts (made within the 24-hour window following the individual's post in the same thread) during the past 7 days in the above test.

The key explanatory variable,  $I(\text{Social Sentiment})_{i,t}$ , is an indicator that takes a value of one if the social sentiment encountered by individual  $i$  on day  $t$ , as defined above, is positive, and zero otherwise. The main dependent variable, Net Buy Dummy $_{i,t+1}$ , is an indicator of whether the individual  $i$  net buys Bitcoin on the next day ( $t + 1$ ). Similar to regression (1), we include control variables such as Bitcoin return and volatility, forum sentiment, and RavenPack news sentiment, which are computed using data over the past 7 days. We also include user-fixed effects and the prior fixed effect. We cluster the standard errors at the user and date levels.

Table 11 presents our findings. As shown in column (1), when social sentiment is positive, investors are 0.37% more likely to net buy the following day. This effect size is economically significant, considering that the average probability of net buying over the next day is only 0.93%. The effect is robust even when we further include control variables and date fixed effects in columns (2) and (3). Overall, sentiment contagion is not a sideshow, as social sentiment significantly predicts individuals' trading in Bitcoin.<sup>11</sup>

Table 12 presents more nuanced evidence regarding the types of users that have a greater influence on others' trading decisions. Following the specifications outlined in Section 4.2, we investigate three pairs of users: naive and legendary, central and non-central, and informed and non-informed users. We then repeat our analysis in regression 2, focusing specifically on the social sentiment from posts published by each type of user. For example, let's consider an investor, denoted as  $i$ , who actively participated in three conversations in the past seven days. In each of these conversations, five different users posted within a 24-hour window after investor  $i$ 's own posts. We group these 15 users based on their types and separately analyze the corresponding social sentiment for each group.

Our findings demonstrate that users' trading decisions are particularly influenced by the social sentiment of specific user groups. First, as indicated in columns (1) and (2), users' trading decisions respond to the social sentiment of naive users, while showing no responsiveness to the social sentiment of legendary users.

---

<sup>11</sup>In Appendix Table IA.9, using a subsample of users who make consecutive posts and also trade Bitcoins around posting on Bitcointalk, we find that investors' belief changes are significantly and positively related to their trading decisions. This is consistent with the findings of Armona et al. (2019) and Giglio et al. (2021).

Second, users' trading decisions display higher sensitivity to the social sentiment expressed by central users, as demonstrated in columns (3) and (4). Third, users' trading decisions are influenced by both informative and non-informative users' social sentiment.

Furthermore, these transaction-based findings serve to validate our primary observations on belief updating. If there exists a strong correlation between belief updates and trading, then the group of users relevant to others' belief updating should significantly overlap with users who exert influence on others' trading decisions. Indeed, this is the case: the response patterns in trading decisions align well with the previously documented belief updating patterns presented in Panel B of Table 7. For example, users' belief updating is generally responsive to the social sentiment expressed by all types of users. Accordingly, we expect their trading decisions to follow a similar pattern, and our analysis in Table 12 confirms this observation.

## **6 Sentiment Contagion and Market Outcomes**

The link between sentiment contagion and individuals' trading decisions in Section 5 provides a foundation for the effect of sentiment contagion at the aggregate level. This section connects sentiment contagion with market outcomes. To analyze market dynamics from a sentiment contagion perspective, we introduce a daily Sentiment Contagion Intensity (SCI) index. This index measures the intensity of sentiment contagion at the market level, based on social interactions among Bitcointalk users in our sample. We demonstrate that SCI contains novel information about future Bitcoin trading volume, volatility, and returns. Importantly, the results in this section also help address concerns about selection issues related to Bitcointalk users. If these users are not representative of Bitcoin investors, we would not expect the SCI (derived from their activities) to have a significant impact on overall market outcomes.

### **6.1 Sentiment Contagion Intensity Index**

We construct the SCI index using a bottom-up approach. First, we identify users whose sentiment changes in the direction of peer sentiment after participating in a conversation on Bitcointalk. The total number of such "infected users" on a given day can measure the intensity of sentiment contagion. To account for the growing number of participants on Bitcointalk over time, we remove the time trend and seasonality by regressing the logarithm of the daily count of infected users on weekday and year-month indicators. The SCI index is then derived as the residual from this regression, normalized to have a mean of zero and a standard deviation of one.

The proposed SCI index shifts the focus from static sentiment levels to the dynamic process of sentiment contagion. Consequently, SCI is expected to capture unique information beyond the average sentiment level. This expectation is preliminarily supported by Appendix Table IA.10, which shows that the contemporaneous correlation between SCI and the average sentiment level in the forum is only 0.6%, a statistically insignificant relationship. Moreover, the information captured by SCI has minimal overlap with the average sentiment level in the news media. Specifically, the correlation between SCI and RavenPack News sentiment is -0.3%, demonstrating that SCI contains distinct information not reflected in news sentiment.

We also investigate the relationship between SCI and several proxies for investor attention. The first proxy, Google search volume (as described in Section 2.4), has a moderate correlation of 23.8% with SCI. The second proxy, the daily number of Bitcoin-related articles published in RavenPack’s news coverage, shows a very low correlation with SCI at 0.4%. Overall, our SCI index captures unique and novel information that is not reflected in traditional proxies for investor sentiment or attention.

## 6.2 Sentiment Contagion and Future Volume

As previously demonstrated in Section 5, sentiment contagion plays a crucial role in the trading decisions of individual investors. Accordingly, we hypothesize a similar link at the aggregate level: the SCI index may serve as a predictor of trading volume in the Bitcoin market.

To test our hypothesis, we conduct a regression analysis using the following equation:

$$\text{Ab Volume}_{t+N} = \beta_0 + \beta_1 \text{SCI}_t + \text{Controls}_t + u_t, \quad (3)$$

where the dependent variable,  $\text{Ab Volume}_{t+N}$ , represents the abnormal dollar trading volume of Bitcoin over the next  $N$  days. Specifically, we calculate the average dollar trading volume (in billions) over the next  $N$  days ( $N = 3, 7, 14$ ) and normalize it by subtracting its average over the past 28 days. The key predictor,  $\text{SCI}_t$ , which measures the intensity of sentiment contagion on day  $t$ . The control variables include events occurring within the past 28 days. To proxy for news arrivals, we calculate the average sentiment score in Ravenpack news during this period. Additionally, we control for Bitcoin volatility and cumulative returns over the same period. Finally, to account for the effect of sentiment levels, we include the average sentiment of all posts on the Bitcointalk forum published during the past 28 days.

Columns (1) to (3) of Table 13 summarize our findings. Overall, the results indicate that SCI significantly and positively predicts abnormal dollar trading volume in subsequent days. For example, in column (1), a one standard deviation increase in SCI corresponds to an additional \$4.76 million in abnormal trading volume on average for the following 3 days. Similarly, column (2) shows that a one standard deviation

increase in SCI results in an additional \$4.35 million in abnormal trading volume per day over the next 7 days. Moreover, the effect of SCI is stronger in the near term than in the longer term. For instance, as the horizon expands to 14 days (column (3)), the average effect of SCI still remains highly significant but the magnitude declines. Notably, the positive relationship between SCI and abnormal trading volume remains robust, even after controlling for news arrivals, recent Bitcoin returns, and overall sentiment levels.

### 6.3 Sentiment Contagion and Future Volatility

Extensive research (Black (1986); De Long et al. (1990)) links investor sentiment to asset price volatility. More recent studies using U.S. stock market data (Da et al. (2014); Antweiler and Frank (2004)) show that sentiment shifts induce noise trading, thereby generating excessive volatility. We extend this line of inquiry by examining it from a social perspective, utilizing our Sentiment Contagion Intensity (SCI) index. A high SCI value reflects intense sentiment propagation through social interactions, which, we hypothesize, triggers more noise trading and increases market volatility.

To formally test this conjecture, we conduct the following regression analysis:

$$rv_{t+N} = \beta_0 + \beta_1 SCI_t + \text{Controls}_t + u_t, \quad (4)$$

where the dependent variable,  $rv_{t+N}$ , represents the average level of realized volatility over the next  $N$  ( $N = 3, 7, 14$ ) days. We compute the realized volatility using the standard deviation of hourly returns within that day. The key explanatory variable of interest,  $SCI_t$ , and the control variables are identical to those specified in equation (3).

Columns (4) to (6) of Table 13 present evidence that SCI positively predicts future volatility in the Bitcoin market. In column (4), a one standard deviation increase in SCI corresponds to a 0.566% increase in the average volatility over the following three days. This increase amounts to 19.32% of the standard deviation of realized volatility (calculated as 0.566% / 2.930%). In columns (5) and (6), we analyze how SCI predicts the average level of realized volatility over the next 7 and 14 days, respectively, and observe a similar pattern. Notably, the effect of SCI is stronger in the short term than in the long term.

Overall, the results in columns (4) to (6) of Table 13 support our hypothesis: as social interactions amplify sentiment contagion, noise trading increases, thereby raising the volatility in the market.

### 6.4 Sentiment Contagion and Price Dynamics

We further demonstrate that sentiment contagion predicts future returns. Intuitively, investors can become either more optimistic (“positively infected”) or more pessimistic (“negatively infected”) through

sentiment contagion, exerting upward or downward noise trading pressure on Bitcoin prices, respectively. We posit that the impact of sentiment contagion on future returns depends on its cumulative net effect. To capture this, we first construct a daily Net SCI index as follows:

$$\text{Net SCI}_t = \text{PosSCI}_t - \text{NegSCI}_t. \quad (5)$$

Here,  $\text{Pos SCI}_t$  represents the number of users “infected” with a positive change in sentiment, while  $\text{Neg SCI}_t$  reflects the number of users “infected” with a negative change in sentiment. Consequently,  $\text{Net SCI}_t$  quantifies the net effect of sentiment contagion on price pressure on day  $t$ .

To evaluate the cumulative net effect, we construct the variable of Cumulative Net  $\text{SCI}_t$  by summing the Net SCI over the past 28 days and normalizing it to have a mean of zero and a standard deviation of one. We hypothesize that the direction of this cumulative net effect predicts future returns. For example, if the cumulative net effect is positive, it indicates that, on average, investors become more optimistic due to social interactions, exerting positive noise trading pressure and pushing Bitcoin prices upward. However, this upward pressure would subsequently lead to price declines as a correction occurs. To formally test this hypothesis, we run the following time-series regression:

$$\text{Market Returns}_{t+N} = \beta_0 + \beta_1 \text{I(Cumulative Net SCI)}_t + \text{Controls}_t + u_t. \quad (6)$$

Here, the dependent variable  $\text{Market Returns}_{t+N}$  represents the cumulative Bitcoin raw returns over the future  $N$ -day window ( $N = 3, 7, 14$ ).  $\text{I(Cumulative Net SCI)}_t$  is a dummy variable that equals one if the cumulative net effect of sentiment contagion is positive over the past 28-day window as of the end of day  $t$ , and zero otherwise. The control variables, which are also constructed using a 28-day window, are identical to those specified in equation (3).

The results in Columns (1) to (3) of Table 14 support our hypothesis. For instance, column (1) shows that when the cumulative net effect is positive, cumulative Bitcoin returns over the next 3 days decrease by 0.845 percentage points, which corresponds to an 11.00% standard deviation decrease (calculated as  $0.845\%/7.684\%$ ) compared to the benchmark case of a negative cumulative net effect. Furthermore, as the time horizon extends, the predictive power of the cumulative net effect remains robust. These findings underscore that sentiment contagion connects to future price dynamics in the Bitcoin market.

Hirshleifer (2020) highlights the error-prone nature of market impact driven by social activities. Building on this insight, we take a step further and hypothesize that sentiment contagion can serve as a predictor of extreme market movements, particularly market crashes. As positive sentiment spreads, it creates sustained buying pressure from successive waves of “infected” investors. This excessive optimism can inflate prices beyond their fundamental value, thereby increasing the likelihood of a sharp correction in the future.

To test this hypothesis formally, we estimate the following regression:

$$\text{Crash}_{t+N} = \beta_0 + \beta_1 \text{I}(\text{Cumulative Net SCI})_t + \text{Controls}_t + u_t. \quad (7)$$

Here, the dependent variable  $\text{Crash}_{t+N}$  is a crash indicator that equals one if at least one daily return falls below the 5th percentile during the future  $N$ -day window ( $N = 3, 7, 14$ ), and zero otherwise. The key explanatory variable and control variables are consistent with those specified in equation (6).

Columns (4) to (6) of Table 14 show that the cumulative Net SCI exhibits significant predictive power for Bitcoin crashes. For example, column (4) indicates that, as the cumulative net effect turns positive, the probability of crashes in the next three days increases by 4.716 percentage points. Moreover, as the future horizon spans, the effect of sentiment contagion on crash further increases: the probability of crashes in the next two weeks further elevates to 14.865 percentage points. Overall, These findings strongly support the predictive power of sentiment contagion for extreme market movements.

## 7 Conclusion

Using unique data on social interactions from Bitcointalk – a prominent online investment platform – we provide direct evidence that sentiment spreads through social networks. Investors systematically adjust their beliefs about Bitcoin prices in the direction of the average peer sentiment encountered during these interactions. Our findings highlight the critical role of social interactions in shaping belief formation.

Moreover, we uncover a non-rational component to the sentiment contagion process. Specifically, while the sentiment expressed in conversations does not predict Bitcoin returns, investors still respond to it, becoming more optimistic (or pessimistic) about Bitcoin following interactions with peers who, on average, express positive (or negative) sentiment. This selective interpretation of social signals in belief updating aligns with confirmation bias. Notably, less central, less informed, and more naive investors are particularly susceptible to peer sentiment in social interactions. Furthermore, contrary to the principles of rational Bayesian updating, the effect of sentiment contagion intensifies when uncertainty about Bitcoin is higher.

Our analysis further demonstrates that social interactions significantly influence both individual investors' trading decisions and broader market outcomes. Consistent with the findings of Giglio et al. (2021), we present robust evidence linking investor beliefs to trading behaviors. By analyzing individual trading records from a subsample of Bitcointalk users, we find a clear and positive correlation between the direction of investors' Bitcoin trades and the peer sentiment they encounter during social interactions.

At the market level, the intensity of sentiment contagion serves as a strong predictor of Bitcoin trading

volume and return volatility. Importantly, we also identify a cumulative net effect of sentiment contagion that significantly predicts future returns. For instance, a positive cumulative net effect indicates that, on average, investors become increasingly optimistic due to social interactions, which exerts positive noise trading pressure on the market, driving Bitcoin prices higher. This optimism, however, often leads to subsequent price reversals. In extreme cases, such cumulative effects from sentiment contagion can even trigger sharp market movements, such as crashes. These findings underscore the error-prone nature of socially driven investment behavior, revealing how sentiment contagion amplifies market inefficiencies and contributes to market fluctuations.

This paper underscores the growing importance of social finance in understanding investor decision-making and market dynamics. The documented phenomenon of sentiment contagion through online conversations offers novel insights into the dynamics of investor beliefs and highlights the broader influence of social interactions on financial outcomes. Online investment communities like Bitcointalk provide a fertile ground for exploring additional questions in the realm of social finance, which we leave for future research.

## References

- Altonji, J. G., Elder, T. E., and Taber, C. R. (2005). Selection on observed and unobserved variables: Assessing the effectiveness of catholic schools. *Journal of political economy*, 113(1):151–184.
- Antweiler, W. and Frank, M. Z. (2004). Is all that talk just noise? the information content of internet stock message boards. *The Journal of Finance*, 59(3):1259–1294.
- Armona, L., Fuster, A., and Zafar, B. (2019). Home price expectations and behaviour: Evidence from a randomized information experiment. *The Review of Economic Studies*, 86(4):1371–1410.
- Baker, S. R., Bloom, N., and Davis, S. J. (2016). Measuring economic policy uncertainty. *The Quarterly Journal of Economics*, 131(4):1593–1636.
- Barber, B. M. and Odean, T. (2000). Trading is hazardous to your wealth: The common stock investment performance of individual investors. *The journal of Finance*, 55(2):773–806.
- Bénabou, R. and Tirole, J. (2016). Mindful economics: The production, consumption, and value of beliefs. *Journal of Economic Perspectives*, 30(3):141–164.
- Black, F. (1986). Noise. *The Journal of Finance*, 41(3):528–543.
- Brown, B., Bakken, J., Ameringer, S., and Mahon, S. (2008). A comprehensive conceptualization of the peer influence process in adolescence. *Understanding Peer Influence In Children and Adolescents*, pages 17–44.
- Brunnermeier, M. K. and Parker, J. A. (2005). Optimal expectations. *American Economic Review*, 95(4):1092–1118.
- Chen, H., De, P., Hu, Y., and Hwang, B.-H. (2014). Wisdom of crowds: The value of stock opinions transmitted through social media. *The review of financial studies*, 27(5):1367–1403.
- Cong, L. W., Li, Y., and Wang, N. (2021). Tokenomics: Dynamic adoption and valuation. *The Review of Financial Studies*, 34(3):1105–1155.
- Cookson, J. A., Engelberg, J. E., and Mullins, W. (2023). Echo chambers. *The Review of Financial Studies*, 36(2):450–500.
- Cookson, J. A., Lu, R., Mullins, W., and Niessner, M. (2024a). The social signal. *Journal of Financial Economics*, 158:103870.
- Cookson, J. A., Mullins, W., and Niessner, M. (2024b). Social media and finance. In *Oxford Research Encyclopedia of Economics and Finance*.
- Da, Z., Engelberg, J., and Gao, P. (2011). In search of attention. *The Journal of Finance*, 66(5):1461–1499.
- Da, Z., Engelberg, J., and Gao, P. (2014). The sum of all fears investor sentiment and asset prices. *The Review of Financial Studies*, 28(1):1–32.
- De Long, J. B., Shleifer, A., Summers, L. H., and Waldmann, R. J. (1990). Noise trader risk in financial markets. *Journal of Political Economy*, 98(4):703–738.
- Feng, L. and Seasholes, M. S. (2004). Correlated trading and location. *The Journal of Finance*, 59(5):2117–2144.
- Giglio, S., Maggiori, M., Stroebel, J., and Utkus, S. (2021). Five facts about beliefs and portfolios. *American Economic Review*, 111(5):1481–1522.
- Griffin, J. M. and Shams, A. (2020). Is bitcoin really untethered? *The Journal of Finance*, 75(4):1913–1964.
- Hagberg, A., Swart, P., and S Chult, D. (2008). Exploring network structure, dynamics, and function using networkx. Technical report, Los Alamos National Lab.(LANL), Los Alamos, NM (United States).
- Han, B., Hirshleifer, D., and Walden, J. (2022). Social transmission bias and investor behavior. *Journal of Financial and Quantitative Analysis*, 57(1):390–412.
- Heimer, R. Z. (2016). Peer pressure: Social interaction and the disposition effect. *The Review of Financial Studies*, 29(11):3177–3209.
- Hirshleifer, D. (2020). Presidential address: Social transmission bias in economics and finance. *The Journal of Finance*, 75(4):1779–1831.
- Hodrick, R. J. (1992). Dividend yields and expected stock returns: Alternative procedures for inference and measurement. *The Review of Financial Studies*, 5(3):357–386.

- Hong, H., Kubik, J. D., and Stein, J. C. (2004). Social interaction and stock-market participation. *The Journal of Finance*, 59(1):137–163.
- Hong, H., Kubrik, J. D., and Stein, J. C. (2005). Thy neighbor's portfolio: Word-of-mouth effects in the holdings and trades of money managers. *The Journal of Finance*, 60(6):2801–2824.
- Huang, S., Hwang, B.-H., and Lou, D. (2021). The rate of communication. *Journal of Financial Economics*, 141(2):533–550.
- Hwang, B.-H. (2023). The impact of word-of-mouth communication on investors' decisions and asset prices. In *Handbook of Financial Decision Making*, pages 171–191. Edward Elgar Publishing.
- Jiang, W. (2017). Have instrumental variables brought us closer to the truth. *Review of Corporate Finance Studies*, 6(2):127–140.
- Kogan, S., Makarov, I., Niessner, M., and Schoar, A. (2024). Are cryptos different? evidence from retail trading. *Journal of Financial Economics*. Forthcoming.
- Kuchler, T., Li, Y., Peng, L., Stroebel, J., and Zhou, D. (2022). Social proximity to capital: Implications for investors and firms. *The Review of Financial Studies*, 35(6):2743–2789.
- Liu, Y., Tsyvinski, A., and Wu, X. (2022). Common risk factors in cryptocurrency. *The Journal of Finance*, 77(2):1133–1177.
- Makarov, I. and Schoar, A. (2019). Price discovery in cryptocurrency markets. In *AEA Papers and Proceedings*, volume 109, pages 97–99. American Economic Association.
- Makarov, I. and Schoar, A. (2020). Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics*, 135(2):293–319.
- Manning, C. D., Surdeanu, M., Bauer, J., Finkel, J. R., Bethard, S., and McClosky, D. (2014). The stanford corenlp natural language processing toolkit. In *Proceedings of 52nd annual meeting of the association for computational linguistics: system demonstrations*, pages 55–60.
- Mian, A. and Sufi, A. (2014). What explains the 2007–2009 drop in employment? *Econometrica*, 82(6):2197–2223.
- Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of general psychology*, 2(2):175–220.
- Oster, E. (2019). Unobservable selection and coefficient stability: Theory and evidence. *Journal of Business & Economic Statistics*, 37(2):187–204.
- Shiller, R. J., Fischer, S., and Friedman, B. M. (1984). Stock prices and social dynamics. *Brookings papers on economic activity*, 1984(2):457–510.
- Shiller, R. J. and Pound, J. (1989). Survey evidence on diffusion of interest and information among investors. *Journal of Economic Behavior & Organization*, 12(1):47–66.
- Simon, D. and Heimer, R. (2012). Facebook finance: How social interaction propagates active investing. In *AFA 2013 San Diego Meetings Paper*.
- Tetlock, P. C. (2007). Giving content to investor sentiment: The role of media in the stock market. *The Journal of Finance*, 62(3):1139–1168.
- Tversky, A. and Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131.

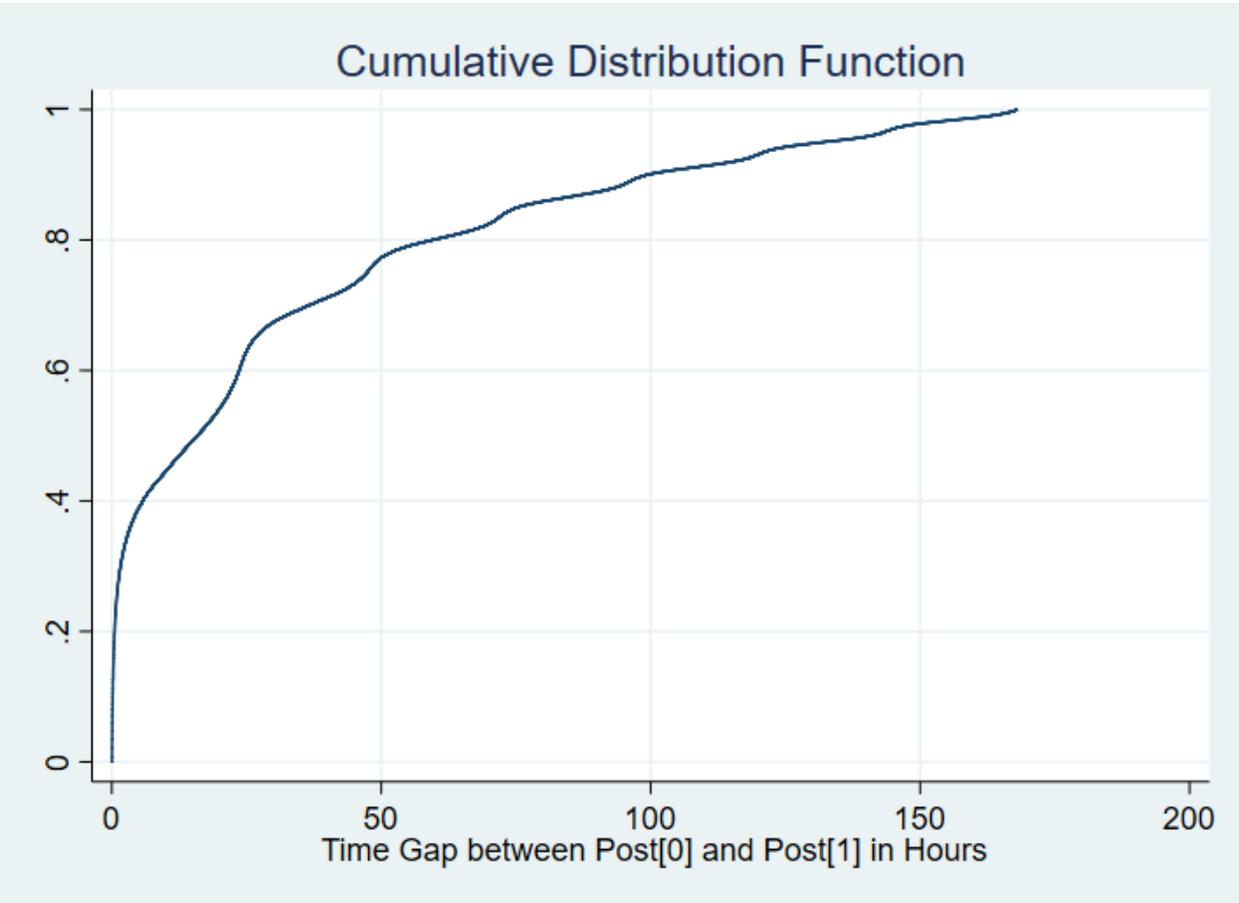


**Figure 1:** Example of Conversation and Illustration of Empirical Strategy

This figure presents an example for a conversation in our sample and our empirical strategy. We measure sentiment changes by calculating the difference between sentiment in post[0] and post[1]. We measure social sentiment by calculating the average sentiment of others' post in conversations.







**Figure 4:** Cumulative Distribution Function of Time Gap

This figure presents the cumulative distribution function of the time gap between two consecutive posts by the same user.

**Table 1:** Summary Statistics

This table tabulates the summary statistics of the key variables. In Panel A, we present statistics of user sentiment in posts aggregated at the user level. To obtain the summary statistics for the average sentiment, we first calculate the mean sentiment of posts published by each user, then report the corresponding statistics based on the whole user population. In Panel B, we present statistics of user sentiment aggregated at the daily level. In Panel C, we present statistics of user sentiment aggregated at the thread level. In Panel D, we present summary statistics of market variables, RavenPack news sentiment and google search volume. Mean Return is the annualized average daily return of Bitcoin. Return Volatility (within Day) is the within-day volatility of Bitcoin based on the hourly Bitcoin return (in percent). Number of Transactions is the total number of Bitcoin transactions within one day. Number of Bitcoins Traded is the total number of Bitcoins being traded within one day. Total Dollar Volume (in millions) is the total dollar trading volume of Bitcoin measured in millions of dollars within one day. RavenPack News Sentiment is the sentiment score for the news on Bitcoin in the RavenPack database. We normalize the sentiment score to be between  $[-1, 1]$ . Google Search (Bitcoin) records the detrended Google Search Index for Bitcoin on a daily basis. Our sample spans from May 1, 2012 to July 30, 2022.

	count	mean	p50	sd	min	max
<b>Panel A: Post Activities at User level</b>						
Average Sentiment	44,356	0.324	0.333	0.512	-1.000	1.000
Standard Deviation of Sentiment	28,077	0.608	0.640	0.297	0.000	1.414
Number of Posts	44,356	15.015	2	72.469	1	3851
<b>Panel B: Post Activities at Daily level</b>						
Average Sentiment	3,737	0.279	0.282	0.118	-0.250	0.833
Standard Deviation of Sentiment	3,736	0.669	0.672	0.054	0.392	0.917
Number of Users	3,737	119.601	85	113.390	1	980
Number of Posts	3,737	178.219	128	160.778	1	1400
<b>Panel C: Post Activities at Thread level</b>						
Average Sentiment	18,799	0.237	0.250	0.324	-1.000	1.000
Standard Deviation of Sentiment	17,191	0.646	0.659	0.189	0.000	1.414
Number of Users	18,799	24.838	11	55.594	1	2589
Number of Posts	18,799	35.428	13	551.522	1	73859
Median Gap in Days between Consecutive Posts	17,214	0.573	0.033	16.568	0.000	1129.115
<b>Panel D: Other Data Sources</b>						
Daily Return(Annualized)	3,743	1.190	0.714	16.166	-177.093	146.518
Return Volatility(Within Day)	3,743	0.037	0.029	0.035	0.000	0.599
Number of Transactions	3,743	16471.536	11376	17268.770	0	181616
Number of Bitcoins Traded	3,743	9358.724	6660.652	9697.557	0	137070.178
Total Dollar Volume (in millions)	3,743	56.499	18.834	97.465	0.000	1372.717
RavenPack News Sentiment	3,743	0.041	0.000	0.330	-0.660	0.660
Google Search (Bitcoin)	3,743	0.031	-0.046	0.417	-0.905	6.595

**Table 2:** Contagion Effect

This table presents the panel regression analysis of sentiment change on social sentiment. The dependent variable is the sentiment change defined as the revision in sentiment between a user's two consecutive posts post[0] and post[1]. The main explanatory variable is the social sentiment defined as the average sentiment of others' posts within the conversation, published between the timestamps of post[0] and post[1]. We control for several contemporaneous variables between post[0] and post[1], such as Bitcoin return and volatility, activities in other conversations on Bitcointalk, and news arrivals according to the RavenPack database. We also include lagged news arrivals to capture delays in response. In all columns, we include a fixed effect for prior sentiment. The *t*-statistics (in parentheses) are based on standard errors clustered by user and day. We multiply each coefficient by 100.

	Sentiment Change				
	(1)	(2)	(3)	(4)	(5)
Social Sentiment	2.838*** (6.81)	2.836*** (6.82)	3.153*** (6.05)	3.096*** (6.02)	3.087*** (6.01)
RavenPack News Sentiment between Post[0] and Post[1]		0.329 (0.41)	-0.271 (-0.32)	-0.195 (-0.23)	0.063 (0.07)
Bitcoin Return			29.443*** (5.76)	29.927*** (5.80)	30.786*** (5.90)
Bitcoin Volatility			15.087 (0.95)	13.439 (0.84)	12.567 (0.78)
Forum Sentiment				-4.659*** (-4.39)	-4.684*** (-4.43)
RavenPack News Sentiment 24 hours before Post[0]					1.593 (1.64)
RavenPack News Sentiment 48 hours before Post[0]					-0.705 (-0.70)
RavenPack News Sentiment 7 days before Post[0]					2.351 (1.21)
Prior FE	YES	YES	YES	YES	YES
User FE	YES	YES	YES	YES	YES
Date FE	YES	YES	YES	YES	YES
Adjusted R-Squared	0.344	0.344	0.345	0.345	0.345
N	212,623	212,623	178,614	175,750	175,750

**Table 3: Placebo Test**

The table displays the results of a placebo test examining the contagion effect. The dependent variable is the sentiment change defined as the revision in sentiment between a user's two consecutive posts post[0] and post[1]. The independent variable is the average sentiment level in  $N$  random conversations ( $N = 1, 3, 5, 10$ ) that the user did not participate in, which occurred between the timestamps of post[0] and post[1]. The control variables remain identical to those in Table 2. In all columns, we include a prior fixed effect. The t-statistics (in parentheses) are based on standard errors clustered by user and day. We multiply each coefficient by 100.

	Sentiment Change			
	(1)	(2)	(3)	(4)
Random Sentiment	0.158 (0.36)	-0.061 (-0.14)	-0.378 (-0.77)	-0.079 (-0.15)
RavenPack News Sentiment between Post[0] and Post[1]	0.216 (0.27)	0.356 (0.42)	0.291 (0.33)	0.903 (0.96)
Bitcoin Return	30.583*** (6.36)	30.916*** (6.13)	31.107*** (6.41)	30.741*** (6.17)
Bitcoin Volatility	9.781 (0.71)	8.801 (0.67)	4.081 (0.29)	11.997 (0.77)
Forum Sentiment	-5.292*** (-5.45)	-7.831*** (-6.24)	-10.699*** (-6.76)	-15.231*** (-7.19)
RavenPack News Sentiment 24 hours before Post[0]	1.232 (1.39)	1.678* (1.80)	1.657* (1.72)	1.520 (1.49)
RavenPack News Sentiment 48 hours before Post[0]	-1.361 (-1.42)	-1.006 (-1.03)	-0.478 (-0.47)	-0.238 (-0.22)
Prior FE	YES	YES	YES	YES
User FE	YES	YES	YES	YES
Date FE	YES	YES	YES	YES
Controls	YES	YES	YES	YES
Adjusted R-Squared	0.345	0.346	0.346	0.346
N	195,635	180,813	169,124	148,284

**Table 4:** Contagion Effect: Oster Test

This table tests the role of omitted and unobservable control variables in driving the relation between social sentiment and sentiment change using the test proposed in [Oster \(2019\)](#). To implement this, we estimate a linear model:

$$\begin{aligned} \text{Senti Change}_{i,j,t_0 \rightarrow t_1} &= \beta_1 \text{Social Sentiment}_{i,j,t_0 \rightarrow t_1} \\ &+ \gamma \text{Control}_{i,t_0 \rightarrow t_1} + \text{Fixed Effects} + u_{i,t_1}. \end{aligned}$$

First, we estimate the model without any control variables, from which we obtain the Social Sentiment coefficient  $\beta_u$  and the coefficient of determination  $R_u^2$ . Then, we include contemporaneous and lagged control variables in our main results, from which we obtain the Social Sentiment coefficient  $\beta_c$  and the coefficient of determination  $R_c^2$  (as reported in the last column of [Table 2](#)). For a given value for the relative degree of selection on observed and unobserved variables ( $\delta$ ) and the R-squared from a hypothetical regression with both observed and unobserved controls ( $R_{max}^2$ ), [Oster \(2019\)](#) derives the bias-adjusted coefficient  $\beta_{adj}$  that is closely approximated by (strictly equal to when  $\delta$  equals one):

$$\beta_{adj} \approx \beta_c - \delta \frac{(\beta_u - \beta_c)(R_{max}^2 - R_c^2)}{R_c^2 - R_u^2}.$$

With this adjusted coefficient  $\beta_{adj}$ , the recommended identified set is the interval between  $\beta_{adj}$  and  $\beta_c$ . In this table, we report the identified set for different combinations of parameters. Additionally, we report whether the identified set rejects the null hypothesis of  $\beta_1 = 0$ .

$R_{max}^2 = \min(2.2R_c^2, 1)$					
$\delta = 1$			$\delta = 2$		
$\beta_{adj}$	Identified Set	Reject Null?	$\beta_{adj}$	Identified Set	Reject Null?
0.083	[0.031, 0.083]	Y	0.177	[0.031, 0.177]	Y

$R_{max}^2 = \min(3R_c^2, 1)$					
$\delta = 1$			$\delta = 2$		
$\beta_{adj}$	Identified Set	Reject Null?	$\beta_{adj}$	Identified Set	Reject Null?
0.105	[0.031, 0.105]	Y	0.281	[0.031, 0.281]	Y

**Table 5:** Contagion Effect: Sentiment Change for Frequent Posters and Infrequent Posters

This table presents subsample analyses of the effect of social sentiment on sentiment change. Specifically, we divide the user sample into frequent posters and infrequent posters. We classify users whose average time elapsed between consecutive pairs of posts is above the population average as infrequent posters, and the remaining users as frequent posters. Column (1) (resp. Column (2)) examines sentiment change for the subsample of frequent (resp. infrequent) posters, using the same regression settings and variable definitions from column (5) in Table 2. We investigate the interactive effect of post frequency in column (3) using the full sample. The t-statistics (in parentheses) are based on standard errors clustered by user and day. We multiply each coefficient by 100.

	Infrequent Posters	Frequent Posters	Interactive Effect
	(1)	(2)	(3)
Social Sentiment	3.503** (2.57)	2.970*** (5.30)	4.095*** (3.34)
Frequent Poster * Social Sentiment			-1.254 (-0.95)
Frequent Poster			0.308 (0.33)
Prior FE	YES	YES	YES
User FE	YES	YES	NO
Date FE	YES	YES	NO
Controls	YES	YES	YES
Adjusted R-Squared	0.343	0.345	0.346
N	31,525	137,437	170,913

**Table 6:** Social Sentiment and Future Returns

This table presents the panel regression analysis of future Bitcoin returns on social sentiment constructed from Bitcointalk posts. The dependent variable is the Bitcoin returns over the next  $K$  hours ( $K = 6, 24, 48, 72,$  and  $168$ ) after  $\text{post}[1]$ . The key independent variable is social sentiment defined as the average sentiment of others' posts between  $\text{post}[0]$  and  $\text{post}[1]$ . Control variables are identical to those in Table 2. In all columns, we include a prior fixed effect. The  $t$ -statistics (in parentheses) are based on standard errors clustered by user and day. We multiply each coefficient by 100.

	(1)	(2)	(3)	(4)	(5)
	6 Hours	24 Hours	48 Hours	72 Hours	168 Hours
Social Sentiment	-0.012 (-0.38)	-0.094** (-2.17)	-0.119*** (-2.69)	-0.111** (-2.31)	-0.114** (-2.57)
RavenPack News Sentiment between Post[0] and Post[1]	-0.733*** (-3.43)	-1.859*** (-5.61)	-1.831*** (-6.76)	-1.369*** (-5.42)	-1.786*** (-5.09)
Bitcoin Return	-21.330*** (-13.37)	-35.052*** (-8.78)	-39.628*** (-7.81)	-41.422*** (-9.48)	-41.485*** (-10.99)
Bitcoin Volatility	17.717 (1.44)	15.433 (0.75)	45.070*** (5.30)	41.646*** (5.36)	37.695*** (3.04)
Forum Sentiment	-0.245*** (-3.82)	-0.537*** (-4.68)	-0.499*** (-4.51)	-0.420*** (-3.68)	-0.418*** (-3.39)
RavenPack News Sentiment 24 hours before Post[0]	-1.566*** (-4.28)	-2.586*** (-5.93)	-2.778*** (-7.26)	-2.223*** (-7.66)	-2.624*** (-6.79)
RavenPack News Sentiment 48 hours before Post[0]	-0.735*** (-2.93)	-1.570*** (-4.10)	-1.030*** (-2.77)	-0.862*** (-2.71)	-1.524*** (-4.67)
Prior FE	YES	YES	YES	YES	YES
User FE	YES	YES	YES	YES	YES
Date FE	YES	YES	YES	YES	YES
Adjusted R-Squared	0.268	0.675	0.824	0.878	0.939
N	175,750	175,750	175,750	175,750	175,748

**Table 7:** Contagion of Sentiment of Different Peer Groups

This table presents how individuals respond to the average sentiment of various types of users with different levels of informedness that they interact with on the Bitcointalk forum. We examine three user features: (i) users labeled as “Legendary” by Bitcointalk; (ii) central users who have participated in more conversations; (iii) informed users whose post sentiment has a high correlation with future Bitcoin returns in 7 days (the correlation exceeds the 25th percentile of the population). In Panel A, we analyze how social sentiment by different types of users predicts future 24-hour Bitcoin returns, using a similar setting in Table 6. In Panel B, we investigate how users revise their beliefs (proxied by post sentiment) following posts from different types of users. We include user and day fixed effects as well as a prior fixed effect in all columns of Panel B. Control variables are identical to those in Table 2. The  $t$ -statistics (shown in parentheses) are based on standard errors that are clustered by user and day. Additionally, we multiply each coefficient by 100 in all panels.

Panel A: Informativeness (future 24-hour returns)						
	Legendary		Central		Informed	
	(1)	(2)	(3)	(4)	(5)	(6)
Social Sentiment by Featured Users	-0.054 (-1.12)		-0.085** (-2.23)		-0.079* (-1.78)	
Social Sentiment by NonFeatured Users		-0.061* (-1.71)		-0.036 (-0.99)		-0.179*** (-2.58)
Prior FE	YES	YES	YES	YES	YES	YES
User FE	YES	YES	YES	YES	YES	YES
Date FE	YES	YES	YES	YES	YES	YES
Controls	YES	YES	YES	YES	YES	YES
Adjusted R-Squared	0.674	0.676	0.673	0.695	0.674	0.676
N	138938	161144	159205	119231	136962	36580

Panel B: Response						
	Legendary		Central		Informed	
	(1)	(2)	(3)	(4)	(5)	(6)
Social Sentiment by Featured Users	0.804* (1.66)		2.897*** (5.55)		2.712*** (4.80)	
Social Sentiment by NonFeatured Users		2.399*** (4.94)		1.925*** (2.84)		4.897*** (3.90)
Prior FE	YES	YES	YES	YES	YES	YES
User FE	YES	YES	YES	YES	YES	YES
Date FE	YES	YES	YES	YES	YES	YES
Controls	YES	YES	YES	YES	YES	YES
Adjusted R-Squared	0.347	0.345	0.345	0.346	0.344	0.338
N	138,938	161,144	173,477	43,701	136,962	36,580

**Table 8:** Effect of Social Sentiment: Heterogeneity across Users

In this table, we explore different users' responses to social sentiment. The dependent variable is the sentiment change defined as the revision in sentiment between a user's two consecutive posts  $\text{post}[0]$  and  $\text{post}[1]$ . In column (1), we examine whether naive users, indicated by a dummy variable that equals one for users with non-legendary status on the Bitcointalk forum, are more susceptible to social sentiment. In column (2), we investigate whether users who have participated in more conversations are less influenced by social sentiment. We measure a user's network centrality based on historical posts, where two users are considered connected if they have participated in the same thread. In column (3), we study whether less informed users, defined as those whose post sentiment has a low correlation with future Bitcoin returns in 7 days (correlation < 25th percentile), are more influenced by social sentiment. In all columns, we include a prior fixed effect. The coefficients are multiplied by 100, and the t-statistics (in parentheses) are based on standard errors clustered by user and day.

	Naive Users	Central Users	Less Informed Users
	(1)	(2)	(3)
User Feature * Social Sentiment	3.585*** (3.46)	-0.047*** (-3.82)	2.664** (1.98)
Social Sentiment	1.125 (1.52)	4.086*** (7.28)	2.105*** (3.49)
User Feature		-0.006 (-0.66)	-0.268 (-0.36)
Prior FE	YES	YES	YES
User FE	YES	YES	YES
Date FE	YES	YES	YES
Controls	YES	YES	YES
Adjusted R-Squared	0.345	0.345	0.343
N	175,750	175,750	139,873

**Table 9:** Time-variation in Sentiment Contagion

In this table, we examine sentiment contagion intensity variations across different time periods. The dependent variable is the sentiment change measured as the revision in sentiment between a user's two consecutive posts  $post[0]$  and  $post[1]$ . In column (1), we investigate the impact of informative days on sentiment contagion by defining informative days based on the number of news arrivals documented in the RavenPack database. If there are no news arrivals on a specific day, we assign a value of 0. We construct a dummy variable called *informative days* that takes the value of 1 on a given day if the total number of news arrivals is higher than the sample median. In column (2), we investigate how information uncertainty affects sentiment contagion. To identify days of high information uncertainty, we use the standard deviation of news sentiment within the day documented in the RavenPack database. We construct a dummy variable called *high information uncertainty*, which takes the value of 1 if this standard deviation is higher than the sample median and 0 otherwise. If there are no news arrivals on a specific day, we assign a value of 0 to the dummy. In column (3), we examine whether Bitcoin return volatility affects sentiment contagion. To do so, we calculate the standard deviation of hourly Bitcoin return within the 24-hour window just prior to the first  $post[0]$ . We construct a dummy variable called *high Bitcoin volatility*, which equals 1 if this standard deviation is higher than the sample mean and 0 otherwise. In all columns, we include a prior fixed effect. The coefficients are multiplied by 100, and the t-statistics (in parentheses) are based on standard errors clustered by user and day.

	Informative Days	High Uncertainty	High Bitcoin Volatility
	(1)	(2)	(3)
Episode Feature * Social Sentiment	0.959 (0.92)	3.851** (2.36)	2.919** (2.50)
Social Sentiment	2.664*** (3.76)	3.613*** (5.42)	2.856*** (3.59)
Episode Feature		-1.059 (-1.10)	-1.501* (-1.96)
Prior FE	YES	YES	YES
User FE	YES	YES	YES
Date FE	YES	YES	YES
Controls	YES	YES	YES
Adjusted R-Squared	0.345	0.322	0.322
N	175,750	143,092	143,092

**Table 10:** Confirmation Bias

This table investigates whether users' responses to social sentiment depends on whether it aligns with their priors. The first two columns utilize the subsample of posts where post[0] has a positive sentiment and examine how changes in sentiment from post[0] to post[1] respond to the positive part and the negative part of social sentiment, denoted by Social sentiment (+) and Social sentiment (-) respectively. In the last two columns, we investigate the response of users with negative priors to positive and negative social sentiment. We include user and day fixed effects in all columns. We report *t*-statistics in parentheses based on standard errors clustered by user and day. Finally, we multiply each coefficient by 100.

	Positive Priors		NonPositive Priors	
	(1)	(2)	(3)	(4)
Social Sentiment	4.279*** (6.06)		3.294*** (4.68)	
Social Sentiment(+)		5.012*** (4.96)		2.764*** (2.74)
Social Sentiment(-)		2.687 (1.61)		4.283*** (2.78)
User FE	YES	YES	YES	YES
Date FE	YES	YES	YES	YES
Controls	YES	YES	YES	YES
Adjusted R-Squared	0.044	0.044	0.249	0.249
N	85,949	85,949	86,237	86,237

**Table 11:** Sentiment Contagion and Individual Trading

This table examines how an individual's trading in Bitcoin depends on the average peer sentiment recently encountered via social interactions (posting on the Bitcointalk forum). The main independent variable is the sign of the average sentiment of others' posts within the conversation (see the definition in Section 5). The dependent variable is an indicator variable for an individual's net purchase the next day. Control variables are identical to those in Table 2. In all columns, we include user fixed effect and a prior fixed effect. We report  $t$ -statistics in parentheses based on standard errors clustered by user and day. We multiply each coefficient by 100.

	Probability of net buying in the next day		
	(1)	(2)	(3)
I(Social Sentiment>0)	0.371*** (3.33)	0.343*** (2.86)	0.293*** (2.61)
Prior FE	YES	YES	YES
User FE	YES	YES	YES
Date FE	NO	NO	YES
Controls	NO	YES	NO
Adjusted R-Squared	0.081	0.081	0.096
N	249,449	249,180	249,449

**Table 12: Heterogeneity in the Effect of Social Sentiment on Trading**

This table presents the heterogeneous effect of peer sentiment on individual trading. The dependent variable is an indicator of whether individuals make a net purchase on day  $t + 1$ . The main explanatory variables are dummy variables corresponding to whether the average sentiment of various groups of peers encountered on day  $t$  is positive. The control variables, listed in Table 2, are measured daily in this table. In all columns, we include user fixed effect and a prior fixed effect. We report  $t$ -statistics in parentheses based on standard errors clustered by user and day. We multiply each coefficient by 100.

	(1)	(2)	(3)	(4)	(5)	(6)
I(Sentiment>0)	0.126 (1.19)	0.123 (1.17)	0.122 (1.14)	0.134 (1.28)	0.086 (0.82)	0.109 (1.03)
I (legendary Exposure Sentiment>0)	0.394** (2.07)					
I (Nonlegendary Exposure Sentiment>0)		0.282** (2.13)				
I (central Exposure Sentiment>0)			0.473** (2.25)			
I (Noncentral Exposure Sentiment>0)				-0.000 (-0.00)		
I (skill Exposure Sentiment>0)					0.580*** (3.86)	
I (Nonskill Exposure Sentiment>0)						0.724*** (3.08)
Adjusted R-Squared	0.096	0.096	0.096	0.096	0.096	0.096
N	249,449	249,449	249,449	249,449	249,449	249,449
User FE	YES	YES	YES	YES	YES	YES
Controls	YES	YES	YES	YES	YES	YES

**Table 13:** Social Sentiment Contagion, Trading Volume, and Market Volatility

This table examines the relationship between social sentiment contagion and future Bitcoin trading volume and volatility. The main explanatory variable is the measure of sentiment contagion intensity (SCI), which is based on the daily number of users who update their beliefs in the same direction as the average peer sentiment in social interactions (see Section 6.1). In columns (1) to (3), the primary dependent variable is the future cumulative dollar trading volume of Bitcoin (in billions) over the next 3 days, 7 days, and 14 days. Trading volume is normalized by subtracting its average over the preceding 28 days. In columns (4) to (6), the main dependent variable is the future volatility of Bitcoin returns over the next 3 days, 7 days, and 14 days. To account for the impact of news arrivals, we include the average sentiment score from RavenPack news over the past 28 days as control variables. Additionally, we incorporate two Bitcoin market variables: Bitcoin volatility and the number of transactions, both measured as cumulative sums over the past 28 days. To capture sentiment in online forums, we calculate the average sentiment of posts on the Bitcointalk forum published during the previous 28 days. We report  $t$ -statistics in parentheses, based on Newey-West corrected standard errors with a 28-day lag. The coefficients are multiplied by 100.

	Abnormal Trading Volume			Return Volatility		
	(1) 3 day	(2) 7 days	(3) 14 days	(4) 3 day	(5) 7 days	(6) 14 days
SCI	0.476*** (3.62)	0.435*** (3.39)	0.325** (2.51)	0.566*** (5.19)	0.446*** (4.51)	0.333*** (3.59)
Bitcoin Volatility in Past 28 Days	-30.182*** (-3.20)	-30.797*** (-3.21)	-30.682*** (-3.12)	33.496*** (7.12)	29.303*** (6.52)	24.912*** (5.26)
RavenPack News Sentiment in Past 28 Days	2.732*** (3.21)	3.073*** (3.30)	3.425*** (3.20)	-0.918 (-1.32)	-1.106* (-1.70)	-1.015* (-1.66)
Forum Average Sentiment in Past 28 Days	-0.254 (-0.17)	-0.013 (-0.01)	0.336 (0.17)	-4.959*** (-3.24)	-4.682*** (-2.98)	-4.495*** (-2.80)
Bitcoin Return in Past 28 Days	1.517** (2.06)	1.397** (2.20)	1.230** (2.26)	2.827*** (2.75)	3.048*** (3.20)	3.077*** (3.53)
Adjusted R-Squared	0.064	0.089	0.104	0.329	0.402	0.439
N	3,734	3,734	3,734	3,734	3,734	3,734

**Table 14:** Social Sentiment Contagion and Market Crashes

This table summarizes the predictive power of social sentiment contagion on Bitcoin returns and market crashes. Social contagion is measured using the cumulative Net SCI, which quantifies the difference between the number of positively affected users and negatively affected users over the past 28 days. The primary explanatory variable is a binary indicator that equals 1 when the cumulative Net SCI is positive and 0 otherwise. Positively affected users are those who increase their sentiment after exposure to positive peer sentiment in conversations, whereas negatively affected users decrease their sentiment after exposure to negative peer sentiment (see Section 6.4). In columns (1) to (3), the dependent variables represent Bitcoin's cumulative raw returns over specified future horizons. In columns (4) to (6), the dependent variable is a crash indicator, which equals 1 if, within a given future horizon, at least one daily return falls below the 5th percentile of all daily returns in the sample period. Control variables are consistent with those used in Table 13. We report  $t$ -statistics in parentheses based on Newey-West corrected standard errors using a lag of 28 days. The coefficients are multiplied by 100.

	Cumulative Returns			Crash (below 5% perc.)		
	(1) 3 day	(2) 7 days	(3) 14 days	(4) 3 day	(5) 7 days	(6) 14 days
I(Cumulative Net SCI >0)	-0.845** (-1.98)	-2.148** (-2.13)	-4.781** (-2.22)	4.716** (1.97)	9.735** (2.31)	14.865*** (2.67)
Bitcoin Volatility in Past 28 Days	15.627 (1.41)	29.264 (1.13)	15.178 (0.36)	186.821*** (3.60)	276.962*** (3.35)	391.176*** (3.63)
RavenPack News Sentiment in Past 28 Days	0.434 (0.35)	1.824 (0.69)	1.120 (0.22)	-1.952 (-0.31)	-3.326 (-0.37)	7.204 (0.64)
Forum Average Sentiment in Past 28 Days	-2.167 (-0.73)	-6.538 (-0.90)	-15.825 (-1.01)	4.629 (0.31)	15.796 (0.60)	22.835 (0.66)
Bitcoin Return in Past 28 Days	2.081 (1.36)	4.557 (1.24)	10.454 (1.38)	6.479** (2.13)	11.136*** (3.04)	10.267** (2.23)
Adjusted R-Squared	0.016	0.032	0.050	0.034	0.057	0.086
N	3,733	3,729	3,722	3,734	3,734	3,734

## INTERNET APPENDIX

### A Assigning Continuous Sentiment Score Using Large Language Models

Stanford CoreNLP serves as a rigorous standard for grammatical sentiment parsing, with the major advantage of effectively avoiding potential look-ahead bias. However, we also employ Large Language Models (LLMs)—specifically the Qwen architecture—as an alternative measure to ensure the robustness of our findings. While Stanford CoreNLP provides a strong baseline, LLMs offer superior contextual interpretation, allowing them to parse complex financial discussions involving sarcasm and conditional statements. This supplementary approach allows us to validate our baseline results.

#### A.1 The Prompt For Extracting Post-Level Sentiment

To construct this alternative measure, we input post-level messages into the Qwen model. If a post consists of multiple sentences, they are grouped together to preserve the full argumentative context.

We utilize a specific prompt designed to enforce a standardized output format and a "Chain-of-Thought" reasoning process. The prompt instructs the model to first verify the relevance of the content before assigning a sentiment value. The exact prompt used is as follows:

*Based on the provided message from forum users, answer the following questions:*

*(1) Does the message contain attitudes regarding Bitcoin's fundamentals, prices, and prospects?*

*Please answer 'Yes' or 'No'.*

*(2) If the answer to question (1) is 'Yes', please assign a score in a continuous range  $[-1, 1]$  to measure the attitude. For example, 1 means extremely optimistic while -1 means extremely pessimistic.*

*(3) Explain the reason.*

This structure ensures that the continuous sentiment score is generated only for relevant financial discussions, minimizing noise in the dataset.

#### A.2 Local Deployment and Implementation Workflow

To ensure data privacy and maximize inference throughput, we implemented a locally deployed LLM (Qwen version: qwen3-14B) rather than relying on external APIs. We utilized the AutoDL cloud computing platform, leasing a server equipped with a single high-performance GPU (e.g., NVIDIA RTX 4090/5090

series). While this hardware configuration theoretically supports models up to 32 billion parameters, we deliberately selected the 14B version to achieve the optimal trade-off between reasoning capability and generation speed.

The deployment workflow was streamlined by leveraging pre-configured environment images, which significantly reduced the overhead of dependency management. We sourced the model weights directly from ModelScope<sup>12</sup> and established the runtime environment using the vLLM library to enable high-efficiency memory management (PagedAttention). The inference engine was initialized via a dedicated shell script to standardize startup parameters, creating a local API endpoint. Finally, we executed the specific sentiment analysis tasks using custom Python scripts designed to interact directly with this locally hosted model, ensuring a stable and controllable analysis pipeline.

### A.3 Illustrative Examples of the Output

The efficacy of this alternative measure is demonstrated by its ability to handle complex linguistic scenarios. Below are examples where the LLM-based method successfully extracts accurate continuous scores, highlighting the functioning of our prompt design.

#### Example 1: Posts with Positive Sentiment

- Post: *The price has been falling for over a year yet more and more VC pours in, more adoption, more transactions, more exchanges, big players, legit exchanges coming onboard. You want to talk about divergence, if there was a graph for the above compared to the price the bullish divergence would be so easy to see. When it takes off again it's really going to knock people's socks off. It is going to be an exponential sling shot. So many people will be trying to get in at the same time there won't be enough BTC available to buy.*
- LLM Output:
  - (1) Relevance: Yes
  - (2) Score: 0.8
  - (3) Reasoning: The message expresses a **strongly optimistic** attitude toward Bitcoin's future prospects despite its current price decline. The user highlights increasing adoption, VC investment, transaction volume, and institutional participation (e.g., "legit exchanges coming onboard") as positive fundamentals, suggesting a divergence between declining prices and growing adoption. They predict a dramatic future price surge ("exponential sling shot") and anticipate a shortage of Bitcoin during the rebound, implying confidence in its long-term value and demand.

---

<sup>12</sup><https://modelscope.cn/models>

While the current price is acknowledged as falling, the focus is on upward momentum and bullish indicators, leading to a high positive score.

### **Example 2: Posts with Negative Sentiment**

- Post: *I do not think Bitcoin replaces Fiat today for these reasons:1. Not trusted. Simply put Bitcoin has value in the market place just not with any of the vendors or people i want to do business with. My local online classifieds returns 9 ads that mention Bitcoin, out of about 100,000 new ads a month. 9. Two years ago it was 15 hits, so trending to zero???*2. *Not portable. Setting up my phone wallet is scary as hell and most certainly not close to mainstream yet.* 3. *Not fast. I cannot complete a person to person transaction in less than 20 minutes.*4. *Not scalable. If I solved 1-3 above, the network would fall on its face.* 5. *Not anonymous. Not a huge deal but nobody is tracking what I buy at garage sales today if I use the good 'ol dollar. Imagine me buying crap from a guy that turns out to be thief and a drug dealer. Nice eternal paper trail i have left behind. I am long Bitcoin, but right now very long. I put this list together thinking about how far we are from true mass adoption. My answer... we are a long, long way out. It really is all about the marriage of software and hardware.... Very complicated problems that will take years and years to iron out. Every mistake leading to losses of millions in stolen coins. This ecosystem is a brutal teacher, if your solution isn't perfect you will lose everything and so will everyone that trusts you. I think it will take 10-20 years of software improvements to solve our current problems, and in the mean time.... I think the price will go lower. Maybe even for years.*
- LLM Output:
  - (1) Relevance: Yes
  - (2) Score: -0.7
  - (3) Reasoning: The message expresses **\*\*pessimistic attitudes\*\*** toward Bitcoin's fundamentals, prices, and prospects. The user highlights critical flaws: lack of trust from vendors, poor portability (complex wallet setup), slow transaction speeds, scalability issues, and insufficient anonymity. While they acknowledge Bitcoin's market value, they argue it is not yet practical for everyday use and predict the price may decline for years due to unresolved technical and adoption challenges. Their emphasis on the "brutal teacher" nature of the ecosystem and the need for "10-20 years of software improvements" underscores a long-term skepticism about Bitcoin's viability as a replacement for fiat, despite their personal investment ("I am long Bitcoin"). The score reflects a strong negative outlook on Bitcoin's current state and near-term prospects.

These examples confirm that our alternative LLM measure effectively parses user intent and provides a high-fidelity, continuous sentiment score that complements our primary Stanford NLP benchmark.

**Table IA.1:** Contagion Effect: Assigning Continuous Sentiment Score using Large Language Models

In this table, we employ Qwen, one of the leading large language models, to assign a continuous sentiment score in the range  $[-1, 1]$  to each post. For a detailed description of the procedure, please refer to the Appendix section A. We then present the panel regression analysis of sentiment change on social sentiment. The dependent variable is the sentiment change defined as the revision in sentiment between a user's two consecutive posts  $\text{post}[0]$  and  $\text{post}[1]$ . The main explanatory variable is the social sentiment defined as the average sentiment of others' posts within the conversation, published between the timestamps of  $\text{post}[0]$  and  $\text{post}[1]$ . Control variables are identical to those in Table 2. In all columns, we include a prior fixed effect. We report t-statistics in parentheses based on standard errors clustered by user and day. We multiply each coefficient by 100.

	Ex-Post Sentiment				
	(1)	(2)	(3)	(4)	(5)
Social Sentiment	2.078*** (4.88)	2.078*** (4.88)	2.802*** (5.46)	2.823*** (5.38)	2.824*** (5.38)
RavenPack News Sentiment between Post[0] and Post[1]		-0.002 (-0.01)	-0.289 (-0.63)	-0.270 (-0.59)	-0.296 (-0.62)
Bitcoin Return			18.047*** (6.73)	18.622*** (6.90)	18.523*** (6.78)
Bitcoin Volatility			-3.698 (-0.39)	-6.414 (-0.68)	-6.346 (-0.67)
Forum Sentiment				-5.587*** (-5.71)	-5.583*** (-5.70)
RavenPack News Sentiment 24 hours before Post[0]					-0.161 (-0.31)
RavenPack News Sentiment 48 hours before Post[0]					0.043 (0.07)
RavenPack News Sentiment 7 days before Post[0]					-0.003 (-0.00)
Adjusted R-Squared	0.310	0.310	0.313	0.313	0.313
N	212,623	212,623	178,614	175,750	175,750
User FE	YES	YES	YES	YES	YES
Day FE	YES	YES	YES	YES	YES
Prior FE	YES	YES	YES	YES	YES

**Table IA.2:** Contagion Effect: Ex-Post Sentiment as Dependent Variable

This table presents the panel regression analysis of ex-post sentiment on social sentiment. The dependent variable, ex-post sentiment, refers to the sentiment in the second post of two consecutive posts (post[0] and post[1]). The main explanatory variable is social sentiment, which is the average sentiment of others' posts within the conversation published between the timestamps of post[0] and post[1]. Control variables are identical to those in Table 2. We also control for the prior sentiment. We report t-statistics in parentheses based on standard errors clustered by investor and day. We multiply each coefficient by 100.

	Ex-Post Sentiment				
	(1)	(2)	(3)	(4)	(5)
Social Sentiment	3.478*** (9.03)	3.474*** (9.03)	4.013*** (8.28)	3.949*** (8.20)	3.940*** (8.18)
Prior Sentiment	-0.498** (-2.20)	-0.498** (-2.20)	-0.495** (-1.97)	-0.541** (-2.15)	-0.544** (-2.16)
RavenPack News Sentiment between Post[0] and Post[1]		0.668 (0.94)	0.296 (0.39)	0.311 (0.41)	0.511 (0.66)
Bitcoin Return			29.835*** (6.54)	30.421*** (6.63)	31.329*** (6.66)
Bitcoin Volatility			-1.623 (-0.10)	-4.302 (-0.26)	-4.900 (-0.30)
Forum Sentiment				-5.862*** (-6.58)	-5.883*** (-6.61)
RavenPack News Sentiment 24 hours before Post[0]					1.554* (1.71)
RavenPack News Sentiment 48 hours before Post[0]					-0.954 (-1.05)
User FE	YES	YES	YES	YES	YES
Date FE	YES	YES	YES	YES	YES
Adjusted R-Squared	0.038	0.038	0.038	0.039	0.039
N	212,623	212,623	178,614	175,750	175,750

**Table IA.3:** Contagion Effect: Max of Sentiment

In this table, we use the max of sentence sentiment to measure sentiment at the post level. We present the panel regression analysis of sentiment change on social sentiment. The dependent variable is the sentiment change defined as the revision in sentiment between a user's two consecutive posts post[0] and post[1]. The main explanatory variable is the social sentiment defined as the average sentiment of others' posts within the conversation, published between the timestamps of post[0] and post[1]. Control variables are identical to those in Table 2. In all columns, we include a prior fixed effect. We report t-statistics in parentheses based on standard errors clustered by user and day. We multiply each coefficient by 100.

	Sentiment Change				
	(1)	(2)	(3)	(4)	(5)
Social Sentiment	2.562*** (6.34)	2.557*** (6.33)	2.814*** (5.61)	2.732*** (5.51)	2.726*** (5.50)
RavenPack News Sentiment between Post[0] and Post[1]		1.111 (1.27)	0.404 (0.43)	0.400 (0.42)	0.463 (0.49)
Bitcoin Return			36.173*** (6.75)	36.908*** (6.84)	37.514*** (6.93)
Bitcoin Volatility			13.072 (0.76)	10.055 (0.58)	9.714 (0.57)
Forum Sentiment				-5.306*** (-5.35)	-5.321*** (-5.37)
RavenPack News Sentiment 24 hours before Post[0]					1.179 (1.08)
RavenPack News Sentiment 48 hours before Post[0]					-1.758 (-1.61)
Prior FE	YES	YES	YES	YES	YES
User FE	YES	YES	YES	NO	NO
Date FE	YES	YES	YES	NO	NO
Adjusted R-Squared	0.385	0.385	0.386	0.386	0.386
N	212,623	212,623	178,614	175,750	175,750

**Table IA.4:** Contagion Effect: Sum of Sentiment

In this table, we use the summation of sentence sentiment to measure sentiment at the post level. We present the panel regression analysis of sentiment change on social sentiment. The dependent variable is the sentiment change defined as the revision in sentiment between a user's two consecutive posts post[0] and post[1]. The main explanatory variable is the social sentiment defined as the average sentiment of others' posts within the conversation, published between the timestamps of post[0] and post[1]. Control variables are identical to those in Table 2. In all columns, we include a prior fixed effect. We report t-statistics in parentheses based on standard errors clustered by user and day. We multiply each coefficient by 100.

	Sentiment Change				
	(1)	(2)	(3)	(4)	(5)
Social Sentiment	2.550*** (5.42)	2.542*** (5.41)	2.849*** (4.99)	2.792*** (4.94)	2.788*** (4.93)
RavenPack News Sentiment between Post[0] and Post[1]		1.804 (1.33)	1.041 (0.71)	1.124 (0.76)	0.927 (0.62)
Bitcoin Return			46.987*** (4.42)	47.556*** (4.44)	47.551*** (4.35)
Bitcoin Volatility			22.847 (1.08)	16.944 (0.79)	17.104 (0.80)
Forum Sentiment				-5.182*** (-4.73)	-5.194*** (-4.75)
RavenPack News Sentiment 24 hours before Post[0]					0.379 (0.23)
RavenPack News Sentiment 48 hours before Post[0]					-3.092** (-2.03)
Prior FE	YES	YES	YES	YES	YES
User FE	YES	YES	YES	NO	NO
Date FE	YES	YES	YES	NO	NO
Adjusted R-Squared	0.310	0.310	0.311	0.311	0.311
N	212,623	212,623	178,614	175,750	175,750

**Table IA.5:** Contagion Effect: Same Conversation

This table presents the panel regression analysis of sentiment change on social sentiment when post[0] and post[1] are published in the same thread within a 24-hours window. The dependent variable is the sentiment change defined as the revision in sentiment between a user's two consecutive posts post[0] and post[1]. The main explanatory variable is the social sentiment defined as the average sentiment of others' posts within the conversation, published between the timestamps of post[0] and post[1]. Control variables are identical to those in Table 2. In all columns, we include a prior fixed effect. We report t-statistics in parentheses based on standard errors clustered by user and day. We multiply each coefficient by 100.

	Sentiment Change				
	(1)	(2)	(3)	(4)	(5)
Social Sentiment	2.165*** (3.14)	2.165*** (3.14)	1.819* (1.93)	1.819* (1.91)	1.821* (1.91)
RavenPack News Sentiment between Post[0] and Post[1]		-1.736 (-0.93)	-1.887 (-0.92)	-1.815 (-0.88)	-1.657 (-0.81)
Bitcoin Return			33.132*** (3.88)	33.386*** (3.84)	33.676*** (3.81)
Bitcoin Volatility			20.181 (0.71)	21.245 (0.75)	20.998 (0.74)
Forum Sentiment				-2.369 (-1.52)	-2.373 (-1.52)
RavenPack News Sentiment 24 hours before Post[0]					0.662 (0.33)
RavenPack News Sentiment 48 hours before Post[0]					0.184 (0.09)
Prior FE	YES	YES	YES	YES	YES
User FE	YES	YES	YES	NO	NO
Date FE	YES	YES	YES	NO	NO
Adjusted R-Squared	0.335	0.335	0.335	0.336	0.335
N	70,388	70,388	54,714	53,283	53,283

**Table IA.6:** Contagion Effect: Sentiment Change over the Next 12, 48, and 72 Hours

This table presents panel regression analysis of sentiment change on social sentiment when post[0] and post[1] are published within a 12-, 48- and 72-hours window. The dependent variable is the sentiment change defined as the revision in sentiment between a user's two consecutive posts post[0] and post[1]. The main explanatory variable is the social sentiment defined as the average sentiment of others' posts within the conversation, published between the timestamps of post[0] and post[1]. Control variables are identical to those in Table 2. In all columns, we include a prior fixed effect. We report t-statistics in parentheses based on standard errors clustered by user and day. We multiply each coefficient by 100.

	12 Hours	48 Hours	72 Hours
	(1)	(2)	(3)
Social Sentiment	2.575*** (4.31)	3.975*** (8.45)	4.235*** (9.17)
RavenPack News Sentiment between Post[0] and Post[1]	0.627 (0.56)	0.220 (0.29)	0.192 (0.28)
Bitcoin Return	28.611*** (4.56)	24.719*** (5.32)	21.591*** (5.37)
Bitcoin Volatility	5.691 (0.31)	14.377 (0.97)	15.457 (1.06)
Forum Sentiment	-3.348*** (-2.99)	-5.744*** (-5.52)	-6.421*** (-6.17)
RavenPack News Sentiment 24 hours before Post[0]	1.211 (0.90)	0.183 (0.26)	0.475 (0.80)
RavenPack News Sentiment 48 hours before Post[0]	-1.692 (-1.11)	-0.806 (-1.13)	-0.200 (-0.33)
Prior FE	YES	YES	YES
User FE	YES	YES	YES
Date FE	YES	YES	YES
Adjusted R-Squared	0.345	0.348	0.350
N	108,372	245,944	283,588

**Table IA.7:** Contagion Effect: More Controls for Lagged Response

This table presents panel regression analysis of sentiment change on social sentiment. For column (1), we adopt the regression settings and variable definitions from column (5) in Table 2. In column (2), we include more controls to account for lagged response to past market dynamics. In column (3), we further account for lagged response to previous posts. In all columns, we include a prior fixed effect. We report t-statistics in parentheses based on standard errors clustered by user and day. We multiply each coefficient by 100.

	(1)	(2)	(3)
Social Sentiment	3.088*** (6.01)	3.066*** (6.03)	3.270*** (6.25)
RavenPack News Sentiment between Post[0] and Post[1]	-0.032 (-0.04)	-0.291 (-0.33)	-0.360 (-0.40)
Bitcoin Return	30.723*** (5.90)	34.684*** (5.52)	33.781*** (5.09)
Bitcoin Volatility	12.925 (0.81)	19.215 (1.28)	10.728 (0.67)
Forum Sentiment	-4.679*** (-4.42)	-4.712*** (-4.55)	-4.440*** (-4.24)
RavenPack News Sentiment 24 hours before Post[0]	1.385 (1.45)	1.295 (1.30)	1.495 (1.48)
RavenPack News Sentiment 48 hours before Post[0]	-1.030 (-1.03)	-0.855 (-0.84)	-0.950 (-0.91)
Bitcoin Return 24 hours before Post[0]		8.698 (1.32)	8.311 (1.23)
Bitcoin Volatility 24 hours before Post[0]		25.187 (0.51)	14.021 (0.29)
Bitcoin Return 48 hours before Post[0]		0.498 (0.08)	0.605 (0.10)
Bitcoin Volatility 48 hours before Post[0]		11.350 (0.24)	16.046 (0.32)
Bitcoin Return 7 Days before Post[0]		-0.329 (-0.06)	1.502 (0.29)
Bitcoin Volatility 7 Days before Post[0]		160.567 (0.95)	149.769 (0.85)
Average Sentiment of Posts before Post[0]			-4.240*** (-4.87)
Prior FE	YES	YES	YES
User FE	YES	YES	YES
Date FE	YES	YES	YES
Adjusted R-Squared	0.345	0.345	0.346
N	175,750	175,749	170,313

**Table IA.8:** Contagion Effect: Sentiment Change for Users with and without Wallet Address

This table presents a subsample analysis of sentiment change on social sentiment. Specifically, we divide the user sample into two groups: users who do not disclose their wallet address and users who voluntarily reveal their wallet address. For each subsample analysis in columns (1) to (2), we adopt the regression settings and variable definitions from column (5) in Table 2. We investigate the interactive effect in column (3). In all columns, we include a prior fixed effect. The t-statistics (in parentheses) are based on standard errors clustered by user and day. We multiply each coefficient by 100.

	Users Without Wallet	Users With Wallet	Interactive Effect
	(1)	(2)	(3)
Social Sentiment	3.088*** (5.31)	3.611*** (2.72)	3.057*** (5.31)
User with Wallet * Social Sentiment			0.187 (0.14)
RavenPack News Sentiment between Post[0] and Post[1]	-0.608 (-0.62)	2.538 (1.14)	-0.032 (-0.04)
Bitcoin Return	30.580*** (5.73)	30.938** (2.08)	30.722*** (5.90)
Bitcoin Volatility	20.176 (1.23)	-54.355 (-1.08)	12.920 (0.81)
Forum Sentiment	-4.457*** (-3.84)	-5.377** (-2.03)	-4.678*** (-4.42)
RavenPack News Sentiment 24 hours before Post[0]	0.569 (0.56)	5.236** (2.27)	1.385 (1.45)
RavenPack News Sentiment 48 hours before Post[0]	-2.111* (-1.86)	3.058 (1.26)	-1.030 (-1.03)
Prior FE	YES	YES	YES
User FE	YES	YES	YES
Date FE	YES	YES	YES
Adjusted R-Squared	0.344	0.341	0.345
N	146,974	28,405	175,750

**Table IA.9:** Sentiment Contagion and Individual Trading: Between Consecutive Pairs

This table presents the regression analysis of individuals' trading decisions on the sentiment change between post[0] and post[1]. The dependent variable is an indicator of whether the individual actively increased her portfolio by at least 0.1 Bitcoins. The main explanatory variable is the sentiment change, defined as the revision in sentiment between a user's two consecutive posts post[0] and post[1]. Control variables are identical to that in Table 2. Additionally, we include the time gap between  $t_0$  and  $t_1$  as an additional control, as trading is more likely to occur over longer time windows. In all columns, we include user and date fixed effects, as well as a prior fixed effect. We report t-statistics in parentheses based on standard errors clustered by user and day. We multiply each coefficient by 100.

	Probability buy in [t0, t1]	Probability buy in [t0, t1], t1-t0>2 hours	Probability buy in in [t1, t1+7days]
Sentiment Change	8.035*** (2.71)	10.229*** (3.35)	0.436* (1.96)
Prior FE	YES	YES	YES
User FE	YES	YES	YES
Date FE	YES	YES	YES
Controls	YES	YES	YES
Adjusted R-Squared	0.238	0.250	0.632
N	203	189	9,284

**Table IA.10:** Correlations between SCI and Other Indicators

This table shows the correlations between SCI indicator and other variables of interest measured at daily frequency. SCI indicator measures the daily number of users who adjust their sentiment to align with the social sentiment in conversations. We use its logarithmic form and eliminate time-fixed effects by regressing on year-month indicators and weekday dummies. Forum Average Sentiment reflects the average level of sentiment in posts published on the forum. RavenPack News Sentiment represents the daily mean sentiment level in news documented in the RavenPack news database. RavenPack News Coverage denotes the number of news articles regarding Bitcoin within a single day. Google Search Bitcoin records the detrended Google Search Index for Bitcoin on a daily basis. Our sample spans from May 1, 2012 to July 30, 2022.

Variables	SCI	Forum Average Sentiment	RavenPack News Sentiment	RavenPack News Coverage	Google Search Bitcoin
SCI	1.000				
Forum Average Sentiment	0.006 (0.696)	1.000			
RavenPack News Sentiment	-0.003 (0.876)	0.347 (0.000)	1.000		
RavenPack News Coverage	0.004 (0.842)	-0.085 (0.000)	-0.082 (0.000)	1.000	
Google Search (Bitcoin)	0.238 (0.000)	-0.032 (0.050)	0.024 (0.265)	0.296 (0.000)	1.000