Social Media Sentiment and Bitcoin Price Dynamics

January 13, 2023

ABSTRACT

Using NLP (natural language processing) and data from Twitter, Reddit, and Stocktwits, this study examines connections among social media sentiment, attention, disagreement, and bitcoin trading activity and returns. We show that 1) higher social media sentiment leads to higher bitcoin returns and trading volume, 2) higher social media attention and disagreement increase bitcoin price volatility, 3) positive changes in sentiment lead to a decrease in volatility, and 4) the magnitude of the impact from different social media varies.

1 Introduction

Despite the ongoing controversy, cryptocurrencies (digital currencies) have undoubtedly been considered by many as a new class of financial instruments, which can be used as alternative currency, asset, and hedging instrument. Cryptocurrencies are based on blockchain technology, a cryptographic and decentralized technology that ensures the digitalization of trust and does not rely on any central authority such as governments or banks. As of October 2021, the total market capitalization of cryptocurrencies reached over $2,500 billion[1] with bitcoin being the most important and by far the largest cryptocurrency in terms of market capitalization (more than $1,000 billion, as of October 2021). Accompanying this sharp increase in cryptocurrency market size has been a remarkable engagement of retail investors and a lack of regulations. Cryptocurrency markets are less regulated than traditional financial markets because, as new global investable instruments traded 24 hours a day over the internet, having a globally legal and synchronized regulation system from all countries is quite difficult. In addition, traditional media is not always interested in timely reporting events involving cryptocurrency, which makes social media a primary source of information. The growing importance of social media in cryptocurrency trading, along with retail investors' considerable amount of time spending on general social media websites (e.g., Twitter) or financial oriented social media websites (e.g., Reddit ad Stock-

[1] https://www.statista.com/statistics/730876/cryptocurrency-maket-value/

Twits), stamps cryptocurrencies as ‘meme’ type security. A meme security refers to the security whose price dynamics are mainly caused by sentiment on social media posts. For example, shares of GameStop stock skyrocketed more than 400% in one week in January 2021 and gained more than 1,600% for the whole month of January 2021. One of the underlying forces of this dramatic movement was amateur traders on “WallStreetBets”, a popular online Reddit forum with more than 10 million active users, to bid up the stock price.

All these social media phenomena lead one to ponder on its role in cryptocurrency trading activities. Today, the question is no longer *whether* social media affects cryptocurrency valuation, but *how* it affects it. With the large amount of data from Twitter, Reddit and Stocktwits, this study first attempts to accurately measure the sentiment embedded in social media by using natural language processing (NLP) models and then investigates the extent to which this sentiment can be used to predict bitcoin trading activities and price dynamics. In addition, we also assess the impact of social media attention and disagreement on the bitcoin market.

Traditional media such as newspapers, online news media and blogs have generally been one-way channels to communicate news and opinions to the general public. In these traditional channels, words and ‘tone’ are carefully chosen by journalists or newspapers’ editors in an effort to present unbiased information. Recent studies analyze firm’s information release and document that both the ‘tone’ and the choice of words in firms’ disclosure documents contain important information and are associated with company performance (Larcker and Zakolyukina (2012), Hobson et al. (2012), and Allee and DeAngelis (2015)). However, with the rise of mobile technologies and online communities, social media has been considered a dynamic two-way channel of information updates (Cade (2018)). The popularity of social media provides both opportunities and challenges for the information environment. On one hand, social media provides an alternative which enables investors to communicate directly their analyses or views without editorial constraints. Due to its network effect, social media can diffuse information more rapidly among targeted groups of audience. Consequently, the proliferation of social media helps to reduce information asymmetry among users and to mitigate adverse market reactions of negative news (Chen et al. (2014), Bartov et al. (2017), Tang (2018), Blankespoor et al. (2014), and Lee et al. (2015)). On the other hand, social media platforms feature social transmission biases (Hirshleifer (2020)) and echo chamber effects. Pedersen (2021) shows that belief that spillovers from social network interactions can lead to excess volumes and volatility. Other empirical research in this area shows that social media widens the reach of false information and exacerbates investors' bias (Demarzo et al. (2003)).

Financial sentiment, broadly defined as the expressed view of a favorable or unfavorable prospect at the basis of an investor's beliefs, has been long posited as a determinant of asset price variation (Keynes (1936)). However, the question of how to accurately measure the sentiment embedded in social media is underexploited. Earlier studies on textual analysis focused solely on the choice and the tone of words by counting the positive or negative words predefined by general-purpose dictionaries (Schrand and Walther (2000), McVay (2006), Larcker and Zakolyukina (2012), and Allee and DeAngelis (2015)). Nevertheless, the positive and negative words defined by general-purpose dictionaries may not be suitable in the financial context. Loughran and McDonald (2011) show that almost three-fourths of the words identified as negative by the widely used Harvard dictionary are words typically not considered negative in a financial context. They further developed an alternative finance-specific sentiment lexicon which since then has been widely used in finance research (Engelberg et al. (2012), Garcia (2013), Chen et al. (2014), among many others[2]) for the analysis of formal financial statements (e.g., annual reports). In the context of social media, even sentiment derived by finance-specific lexicon might be biased for several reasons. First, posts on social media often use non-standard informal English language (Liu et al. (2012)). Second, social media posts are often written in a social setting, captures communications among a group of people with common interests (Park et al. (2015)). Third, languages used in social media posts present individuals' own views about the world (Back et al. (2010)). Dictionary-based measures may not be able to correctly identify the financial sentiment contained in these statements. In this paper, we contribute to the literature by filling this gap. Specifically, our study adopts a cutting-edge

[2]see Loughran and McDonald (2016) for surveys.

NLP model that is trained to measure the textual sentiment specifically in the context of finance. The model is far superior to dictionary-based methods in text understanding because it can capture the context and the order of words by treating a text as a sequence of words.

Our paper contributes to the literature in three important empirical dimensions. First of all, to the best of our knowledge, we are among the first to use an NLP model to measure financially-oriented sentiment embedded in social media. Several novel results emerge from our results. Second, we show that social media sentiment exhibits a Granger causality to future bitcoin returns and trading volumes, but not to future volatility. Social media sentiment instead has a contemporaneous (same-day) effect on volatility. More specifically, positive changes in sentiment lead on average to a decrease in volatility during the day. We also find that the impact of sentiment on bitcoin trading is different among the three social media platforms that we consider. Third, in addition to sentiment, our study also shows that a rise in attention and disagreement increases uncertainty by raising volatility and skewness, to a lesser extent.

The rest of the paper proceeds as follows. In Section 2, we describe the data that we use in our empirical tests. In Section 3, we discuss the construction of social media sentiment measures. In Section 4, we present and discuss results of empirical tests. We offer robustness checks in Section 5. Section 6 concludes.

2 Data

In this study we use six different datasets: 1) social media textual data from Twitter, Reddit and StockTwits, 2) articles released in traditional media outlets (e.g., Wall Street Journal), 3) daily Bitcoin Price from Coinmarketcap.com, 4) intraday bitcoin transaction data from Kaiko, 5) Google search data series, and 6) daily financial index data, obtained from Yahoo Finance.

To obtain the sentiment embedded in social media, we first scrape bitcoin-related social textual data via the Application Programming Interface (API) provided by Twitter, Reddit and StockTwits for the period between January 2017 and December 2020. Reddit and StockTwits typically feature discussions from more financially-savvy users and offer an advantage to extract the sentiment of cryptocurrency traders, which may ultimately have an impact on bitcoin short-term returns.[3] On the other hand, Twitter offers a relatively "noisy" sentiment because postings in Twitter also contain general news. For Twitter and Reddit, the posting messages are scraped with the keyword "Bitcoin". In StockTwits, one can filter a cryptocurrency with a hashtag that ends with ".X" (e.g., $BTC.X for Bitcoin). We use this convention to download all postings related to Bitcoin in StockTwits during our sample period. The scraped postings of these social media include posted messages, dates and timestamps. For our sample period, our final bitcoin-related textual dataset includes 28.7 millions messages from Twitters, 6.57

[3]See Betzer and Harries (2021), Hu et al. (2021), Diangson (2021), Agrawal et al. (2018), and Awais and Yang (2021)

millions from Reddits, and 1.14 millions from StockTwits. With this dataset, we further measure the sentiment of every post/tweet using an NLP learning model, which is documented in Section 3. The sentiment of a post/tweet is a continuous numeric value between -1 (negative) and 1 (positive). Daily sentiment is then defined as the average sentiment and disagreement as the standard deviation of the sentiment of every messages posted during a given day. Attention is proxied by the number of postings during the day as in Da et al. (2011).

In order to test whether social media sentiment has additional impact on bitcoin beyond traditional media sentiment, we also compute sentiments of bitcoin embedded in Wall Street Journal (WSJ), Dow Jones Newswires (DJN), and Reuters. The articles are collected via Factiva database. Specifically, to avoid the articles about bitcoin related companies instead of bitcoin, we take bitcoin as our key word and choose cryptocurrency market as the main subject as our search criteria. Finally, we obtain a total of 1,450 articles published in the three traditional media. The numbers of articles for DJN, Reuters, and WSJ are 614, 577, and 259, respectively. On average, there are at lease two articles per day that can be used to compute the sentiment of traditional media during our sample period. As for sentiment computation, we apply the same NLP algorithm as for social media postings.

Our daily bitcoin price and volume data are from Coinmarketcap.com, which is a leading source of cryptocurrency data. It collects and aggregates information from over 200 major exchanges and provides daily data on open, close, high, low prices, and volume. For each cryptocurrency, Coinmarketcap.com calculates its volume-weighted price of all price reported at each exchange. To conduct our analysis of realized volatility, we also use the intraday-level data. The intraday transaction data used in this paper are from the leading cryptocurrency market data provider Kaiko. Its raw cryptocurrency data covers 20,000+ pairs across worldwide exchanges. Our dataset is at the tick-by-tick level, including unique trade id, exchange code, currency pairs, prices, volumes, trade directions, and timestamps, for all exchanges where Bitcoin is traded.

Google search data series for the word ”Bitcoin” are downloaded from Google. We further reconstruct Google trends daily data as in Liu and Tsyvinski (2020). The market daily index series are from Yahoo finance. The indexes used to capture financial market dynamics include SP&500, MSCI Global, Gold Index, USD Index, VIX, and U.S. Treasury bond yield.

3 Social Media Sentiment

3.1 Background

With the growing availability of digital textual data and computing technology, the measurement of financial sentiment embedded in texts has received increased research interest. Some pioneering papers (Tetlock, 2007; Kothari, Li and Short, 2009) popularized the simple dictionary-based approach (i.e., counting within a text the presence of positive and negative words predefined by a sentiment dictionary). Later, Loughran and McDonald (2011) pointed out that the general-purpose dictionaries commonly used by previous researchers often misclassify words in the financial context. They curated a finance-specific sentiment lexicon, which has since been broadly adopted by other researchers in the domain of finance. Besides the dictionary-based approach, classical machine learning methods such as naïve Bayes have also been widely explored (Antweiler and Frank, 2004; Das and Chen, 2007; Huang, Zang and Zheng, 2014). Those methods are usually supervised, training statistical models that learn from examples of texts with sentiment labeled by human experts. As shown in the aforementioned papers, finance researchers traditionally rely on the bag-of-words methods, which treat a text as a collection of independent words that can be represented by a vector of word counts. Due to their simplicity, the bag-of-words methods lack the ability to capture the context and the order of words which are crucial for interpreting the semantics of a text.

Recently, the development of natural language processing techniques has introduced more sophisticated models that are capable of recognizing the sequential nature of text and preserving the dependencies between words. The past decade has witnessed rapid advances in the field of NLP (Mikolov, Chen, Corrado and Dean, 2013; Bahdanau, Cho and Bengio, 2014; Cho, van Merrienboer, Gulcehre, Bahdanau, Bougares, Schwenk and Bengio, 2014; Pennington, Socher and Manning, 2014), with the help of deep learning, a subset of machine learning that features deep neural network models capable of tackling complex unstructured data such as texts and images. Specifically, revolutionary breakthroughs have been achieved recently by the novel transformer-based language models Vaswani, Shazeer, Parmar, Uszkoreit, Jones, Gomez, Kaiser and Polosukhin (2017) such as GPT (Generative Pretrained Transformer) (Radford, Narasimhan, Salimans and Sutskever, 2018; Brown, Mann, Ryder, Subbiah, Kaplan, Dhariwal, Neelakantan, Shyam, Sastry and Askell, 2020) and BERT (Bidirectional Encoder Representations from Transformers) (Devlin, Chang, Lee and Toutanova, 2019), which successfully bring the model performance on many NLP tasks to the human level. Those transformer models are gigantic in size, often have hundreds of millions of parameters. A such model is first pretrained on large unlabeled text datasets like the whole Wikipedia corpus, so that it can encode abundant linguistic knowledge. Then it only requires a small amount of labeled data to be fine-tuned on specific tasks such as sentiment analysis, due to its ability to transfer the knowledge it has learned from the unlabeled corpus to the downstream tasks. Incentivized by the breakthroughs in the NLP, some researchers begin to explore those cutting-edge models for their application in finance. (Araci, 2019; Yang, Uy and Huang, 2020) both show that by pretraining the BERT model on finance specific corpora and then fine-tuning it on sentiment analysis, the model can achieve state-of-the-art performance for various financial sentiment analysis datasets.

In this paper, we develop our own transformer model named "FinRoBERTa" to compute the financial sentiment measurement on the social media postings associated with bitcoin. The backbone architecture of the FinRoBERTa is the cutting-edge RoBERTa (Robustly optimized BERT approach) model (Liu, Ott, Goyal, Du, Joshi, Chen, Levy, Lewis, Zettlemoyer and Stoyanov, 2019), an improved derivative of the BERT model. We pretrain the model on finance domain corpora and fine-tune it on financial sentiment data. Our model achieves state-of-the-art performance for financial sentiment analysis on the test set.

3.2 FinRoBERTa Financial Sentiment Model

Model Pretraining

The purpose of pretraining a language model is to leverage large text corpora as self-labeled data to teach the model language knowledge. The RoBERTa model that we adopt in this paper uses the Masked Language Model (MLM) technique in which we randomly mask some of the words from the input text, and let the model try to predict the masked word based on its context (Devlin et al., 2019).

The original base version of the RoBERTa model has 12 layers of the transformer neural network modules, comprising 110 million parameters. It's pretrained on 160GB of general English-language corpora including books, news, online texts, etc., and takes days with over a thousand Nvidia V100-32GB GPUs, a massive amount of computing resources (Liu et al., 2019). Through pretraining, the model learns rich general language knowledge.

However, the word distribution of financial corpora can be quite distinct from that of the general corpora because the financial domain uses a lot of its own techincal jargon. Researchers in different specialized domains have reported that a transformer model pretrained on domain corpora can outperform the generic model on domain-specific tasks (Huang, Altosaar and Ranganath, 2019; Lee, Yoon, Kim, Kim, Kim, So and Kang, 2019). In this regard, we pretrain the FinRoBERTa model from scratch on a 2.6GB English corpora of 2.5 million financial news collected from Factiva database. To reduce the computational needs, we adopt a smaller version of the RoBERTa model architecture with 6 transformer layers totaling 57M parameters, and decrease the vocabulary size from the original 50K to 30K.

The model was implemented using the Huggingface Transformers python library (Wolf, Chaumond, Debut, Sanh, Delangue, Moi, Cistac, Funtowicz, Davison and Shleifer). We pretrain the model for 4 epochs (cycles over the whole dataset) on a server equipped with 4 Nvidia V100-16GB GPUs.[4] The pretraining took around 80 hours.

After pretrained on the financial corpora, the FinRoBERTa model can demonstrate a grasp of financial language knowledge when fulfilling the MLM task, as shown in the example in Table 1.

[Insert Table 1 here]

It is also worth noticing that the vocabulary the model learns from the

[4]The computing resource is supported by Calcul Québec (www.calculquebec.ca) and Compute Canada (www.computecanada.ca).

financial corpora features many common financial technical terms that are not captured by the generic RoBERTa vocabulary, such as “IPO”, “EPS”, “ROE”, “EBITDA”, “GAAP”, “CFA”, “WSJ”, etc. This added awareness of finance terminology would almost certainly contribute to its superior performance on finance specific tasks.

Model Fine-tuning

Fine-tuning is the process of further training the pretrained model on the labeled data of a target task such as sentiment analysis, so that the model learns to solve the specific task. Since the pretrained language model has already encoded abundant language knowledge, it can solve the end task much better given a limited amount of labeled training data, compared to traditional machine learning models that are trained merely on those labeled data.

The fine-tuning is performed with the financial phrase bank (FPB) of (Malo, Sinha, Korhonen, Wallenius and Takala, 2014). It contains around 5,000 sentences randomly selected from financial news. Each sentence is manually labeled with the financial sentiment as either negative, neutral, or positive, independently by 5 to 8 annotators with adequate background in finance and business. The data is divided into 4 subsets that each contain sentences meeting a certain agreement level, i.e., the percentage of annotators agreeing on the same label for a sentence, namely 100%, >75%, >66% and >50%. Crucially, the financial sentiment here is different from ordinary sentiment: it's defined as the potential impact of a new information on future financial events from an investor's point of view.

We follow the same fine-tuning method as in the RoBERTa and BERT papers, by appending a 2-layer neural network for classification to the Fin-RoBERTa model, that takes the sentence embedding from the transformer layers as input and outputs the predicted probabilities for the 3 sentiment classes: Negative, Neutral, and Positive. The training criteria is the cross-entropy loss which measures the divergence of the predicted class probabilities from the true class. The loss of each class is adjusted by a weight of $1/\sqrt{\%\ of\ the\ class\ in\ the\ data}$ to alleviate the class imbalance issue.

Considering that a sentence without a clear majority agreement by people should be very vague in its sentiment, it's hard to justify using it as golden standard for the model. So, we chose to fine-tune our model only on the FPB subdatasets with 100%, >75%, and >66% agreement separately (FPB-100, FPB-75, FPB-66). We split each data into 3 sets: 60% for training, 20% for validation, and 20% for test.

The test results on the different datasets are shown in table 2. The performance is in par with the state-of-the-art performance by similar finance specific transformer models of even larger size (Araci, 2019; Yang et al., 2020), and it exceeds the traditional dictionary approach (Loughran and McDonald, 2011) and classical machine learning approach (Malo et al., 2014) by a large margin.

[Insert Table 2 here]

3.3 Sentiment Measurement on Bitcoin Social Media Postings

We apply the fine-tuned FinRoBERTa model to measure the financial sentiment of the social media postings described in the last section. In view of the tradeoff between label quality and data quantity, we choose to use the model fine-tuned on the FPB-75 dataset. We first clean the texts to get rid of the noises such as web address. Then we input each text to the fine-tuned FinRoBERTa model to predict its financial sentiment. The output sentiment score is between -1 and +1, calculated as the predicted probability of being positive minus that of being negative, so that a more negative (positive) score means more negative (positive) sentiment. A score of 0 is interpreted as neutral since it has an equal probability of being positive as being negative. The measured financial sentiment distributions of postings from different social media platforms are shown in table 4.

To further validate our measurement, we manually compare the financial sentiment measured by our FinRoBERTa model against the general sentiment measured by a conventional NLP model and the financial sentiment measured with the classical Loughran-McDonald dictionary. For the general sentiment, we apply TextBlob[5], a popular NLP library that uses an expert-crafted English sentiment lexicon and linguistic rules to measure sentiment in a text. By manually examining the sample results, we confirm that our

[5]https://textblob.readthedocs.io/en/dev/

financial sentiment captures well a text's financial implication from an investor's perspective, whereas the general sentiment and the dictionary-based financial sentiment often fail badly. Table 3 shows 10 representative examples that compare the three sentiment measures.

[Insert Table 3 here]

4 Empirical Results

4.1 Summary Statistics

Table 4 presents the summary statistics for the variables used in our study. Panel A shows the key statistics of bitcoin daily returns, volumes (in million of bitcoins, denoted by Vol), sentiment (denoted by Sent), number of postings (denoted by Nb), and disagreement (denoted by Dis). For sentiment-related variables, we present the statistics for the three social media sources (i.e., Twitter, Reddit, and StockTwits). Bitcoin has a daily average return of 0.07% with a skewness of -1.48 and kurtosis of 21.08, suggesting an overall increase during the analyzed period but accompanied by more negative observations and a relatively large number of extreme values. Figure 1 plots daily realized variance and kurtosis from January 2018 to January 2021. The realized vairance and kurtosis were at their peak in March 2020. The corresponding trading volume during this period has a mean 18.62 million with a standard deviation of 13.80 m which implies that 68% of observations fall into the large interval between 4.82 and 32.42 billion dollars. The average sentiment for each of the three social media is all slightly positive (0.04 for Twitter, 0.02 for Reddit, and 0.03 for StockTwits) during the sample period, in line with an overall positive sentiment during the sample period. The number of postings, a proxy for attention, varies with social media. Twitter, the largest social media in the world, has a mean of 19.45 thousand bitcoin-related postings per day. Reddit (Finance subreddit) and StockTwits are more financially oriented social media and contain, on average, 4.1 and 1.0 thousand postings per day, respectively. Another sentiment-related measure, disagreement, is around 0.27 and remains stable among three social media.

[Insert Table 4 here]

Table 5 presents the correlation matrix for bitcoin returns, social media sentiments (raw and orthogonal), and traditional media sentiments. The orthogonal sentiments are the residuals from the regression of sentiments on lagged bitcoin returns. On average, the correlations between return and the three social media are around 0.45, which is quite similar for orthogonal sentiment but much higher than the correlation between return and traditional media (0.13).

[Insert Table 5 here]

4.2 The Impact of Social Media on Bitcoin Volume and Return

We first look at how social media affects bitcoin daily trading dynamics such as returns and volume. We estimate the following VAR model with daily bitcoin returns and various social media measures (sentiment, attention and disagreement).

$$x_t = c_x + \sum_{\tau=1}^{2} \alpha_{x,\tau} x_{t-\tau} + \sum_{\tau=1}^{2} \alpha_{y,\tau} y_{t-\tau} + \alpha_z Z_{t-1} + \varepsilon_{x,t}, \tag{4.1}$$

$$y_t = c_y + \sum_{\tau=1}^{2} \beta_{x,\tau} x_{t-\tau} + \sum_{\tau=1}^{2} \beta_{y,\tau} y_{t-\tau} + \beta_z Z_{t-1} + \varepsilon_{y,t}, \tag{4.2}$$

where x_t and y_t are variables of interest on day t, which include bitcoin daily returns (Ret_t), bitcoin daily trading volume (Vol_t), social media sentiment (i.e., Twitter $Sent_t^T$, Reddit $Sent_t^R$, or Stocktwits $Sent_t^S$), social media attention (Nb_t), and disagreement (Dis_t). Z_{t-1} represents control variables (e.g., traditional media sentiment ($Sent_{t-1}^{trad}$)).

The results are reported in Table 6. Coefficients of the sentiment of day $t-1$ on returns on day t are statistically significant and positive for Twitter and StockTwits, suggesting that a higher social media sentiment on a given day can lead to a positive bitcoin return the next day. The opposite causality, i.e. a positive bitcoin return on day $t-1$ also leads to a higher social media sentiment on day t given that the coefficients of bitcoin returns are also statistically significant and positive. However, this Granger causal relationship of bitcoin returns on sentiment is smaller than that of social media sentiment on bitcoin returns.

[Insert Table 6 here]

We further apply the VAR model to daily bitcoin trading volume and the three social media sentiment measures. The results show that a positive sentiment of Twitter and Reddit on bitcoin can lead to a significant increase in bitcoin trading volume. However, the opposite is not true in a statistically significant way. Combined with the results of bitcoin returns and sentiment, we conclude that a positive sentiment results in a stronger buy intention and then leads to higher returns. The results further indicate that attention, measured by the number of postings, has a time-varying impact on trading volume. Specifically, more attention in social media can lead to an increase in trading volume next day. However, this increase will be offset by a decrease in trading volume in two days.

Finally, social media disagreement of Twitter and Reddit has a significant net negative impact on bitcoin trading volume, while social media disagreement of StockTwits has a significent postive impact on bitcoin trading volume. The intuition is that when there is more difference in opinion, investors tend to trade less the cryptocurrency. On the other hand, a higher trading volume can lead to different levels of disagreement for Twitter and Reddit. Recall that the profiles of Twitter and Reddit users are more general than those of StockTwits. Our results suggest that the opinion divergence is more persistent in Twitter and Reddit than that in StockTwits which contains more financially oriented users.

It is worth noting that the sentiment of traditional media does not affect the bitcoin return and trading volume in general. When the sentiment of traditional media is used as a control variable in various VAR models as reported in Table 6, its coefficients are not significant except for one case. In the VAR model of bitcoin return and the sentiment of traditional media, the lagged traditional media sentiment has insignificant coefficient, while the past bitcoin returns have positive and significant impact on the traditional media sentiment.

4.3 Determinants of Bitcoin Price Volatility and Higher Moments

We now turn our attention to the impact of social media on bitcoin price volatility, skewness and kurtosis. We consider the following OLS model:

$$MoM_t = \beta_0 + \beta_1 \times MoM_{t-1} + \beta_2 \times SocMe_{t-1} + \beta_3 \times Ret_{t-1} + \beta_z Z_t + \epsilon_t, \tag{4.3}$$

$$RV_t = \sum_{i=1}^{n} r_{i,t}^2, \tag{4.4}$$

$$Skew_t = \frac{\sum_{i=1}^{n} r_{i,t}^3}{(n-1) \times \sigma_t^3}, \tag{4.5}$$

$$Kurt_t = \frac{n(n+1)}{(n-1)(n-2)(n-3)} \frac{\sum_1^N r_{i,t}^4}{\sigma_t^4}, \tag{4.6}$$

where MoM_t stands for bitcoin daily realised variance (RV_t), skewness ($skew_t$), or kurtosis ($kurt_t$) on day t, all three defined as the median over the nine most active bitcoin exchanges.[6] $SocMe_t$ corresponds to the social media related variables (i.e., sentiment, attention, and disagreement) and Ret_{t-1} is the lagged bitcoin daily return. $r_{i,t}$ is the i-th 5-min bitcoin return on day t and $n = 288$ is the number of 5-min interval during the day. Z_t represents control variables (e.g., traditional media sentiment ($Sent_t^{trad}$))

Table 7 shows that, without including the lagged bitcoin return as one of control variables, social media sentiment has a significant negative impact on realized volatility. However, when controlling with lagged bitcoin returns, the social media sentiment no longer has a significant impact on bitcoin future price volatility. In these three cases, lagged bitcoin return has a significant negative impact, providing evidence on the phenomenon known as the leverage effect in asset pricing literature (Bollerslev et al. (2006), Carr and Wu (2017), and among others).

[Insert Table 7 here]

Table 8 provides mixed evidence, after controlling for lagged returns, that social media sentiment has little or no impact on daily return skewness but a significant impact on daily kurtosis. Given that sentiment has a significant positive impact on bitcoin returns, the result implies that positive sentiment

[6] These nine exchanges are Bibox, BeQuant, BitForex, Bit-Z, Binance, EXX, Huobi, OkEX, and ZB.

is likely to cause more extreme bitcoin returns, but not a mild asymmetry of returns.

[Insert Table 8 here]

By putting all three social media variables together, Table 9 confirms that sentiment does not affect future volatility, but attention and disagreement do. Regarding the intraday bitcoin return skewness, only the coefficients of sentiment and attention from StockTwits are positively significant. Further, disagreement of Twitter and Reddit have a positive significant impact on bitcoin return skewness, suggesting that when there is a divergence in social media sentiments, it is more likely to observe more positive intraday bitcoin returns. Finally, the bitcoin intraday returns' kurtosis is related to sentiment and disagreement, but not the attention. More specifically, a more positive (negative) sentiment from Twitter and Stocktwits result in an increase in probability of extreme positive (negative) returns. Also, our results suggest that disagreement from Twitter and Reddit is also important factor to drive more extreme observations.

[Insert Table 9 here]

Table 10 shows that it is not the lagged social media sentiment, but contemporaneous social media sentiment, that affects bitcoin price volatility. Our results show that the coefficient of sentiment variation ($\Delta S_t = S_t - S_{t-1}$) at day t has a significant negative impact on volatility at the same day, suggesting that a rise in positive sentiment can reduce bitcoin price volatility during the same day, even when the lagged returns variable is included as a control. The results in Table 10 also confirm the positive relation between sentiment variation and bitcoin intraday return skewness and the positive relation between lagged social media sentiment variation and bitcoin intraday return kurtosis.

[Insert Table 10 here]

5 Robustness check

5.1 VAR Model with Additional Control Variables

In Table 11, we revisit the previous VAR models with common financial indices as controlled variables:

$$x_t = c_x + \sum_{\tau=1}^{2} \alpha_{x,\tau} x_{t-\tau} + \sum_{\tau=1}^{2} \alpha_{y,\tau} y_{t-\tau} + \sum_{\tau=1}^{2} \alpha_{z,\tau} z_{t-\tau} + \varepsilon_{x,t}, \quad (5.1)$$

$$y_t = c_y + \sum_{\tau=1}^{2} \beta_{x,\tau} x_{t-\tau} + \sum_{\tau=1}^{2} \beta_{y,\tau} y_{t-\tau} + \sum_{\tau=1}^{2} \beta_{z,\tau} z_{t-\tau} + \varepsilon_{y,t}, \quad (5.2)$$

where z_t is control variable for VAR model and x_t and y_t have the same definitions as in equations 4.1 and 4.2.

We used the following controlled variables in the above regressions: lagged traditional media sentiment, MSCI World Index, US dollar index (DXY), gold prices, Invesco DB Commodity Index, Dow Jones Commodity Index (DJCI), crude oil prices, SPDR S&P 500 ETF, VIX volatility index and Yield of U.S. 10-year treasury note (TNXT). Table 11 confirms the results of the relationship between return, volume, and social media related variables. Specifically, the results show that 1) a higher social media sentiment can lead to a positive bitcoin return next day, 2) a positive sentiment on bitcoin can lead to an increase in bitcoin trading volume, however, the opposite is not always true, 3) more attention in social media can lead to an increase in trading volume next day, however, this increase will be offset by a decrease in trading volume in two days, 4) social media disagreement has a significant net negative impact on bitcoin trading volume.

[Insert Table 11 here]

5.2 Principal Components of Social Media-related Variables

$$MoM_t = \beta_0 + \beta_1 \times MoM_{t-1} + \beta_2 \times PC_{t-1}^{SocMe} + \beta_3 \times Ctrol_{t-1} + \epsilon_t, \quad (5.3)$$

where PC_t^{SocMe} relates to the principal components of the corresponding social media related variables. MoM_t, $SocMe_t$, and $Ctrol_t$ have the same definitions as in equation 4.3

Using principal components to capture information embedded in three social media- related variables, Table 12 indicates that lagged social media sentiment do not have impact on bitcoin price volatility, however, social media attention does have.

[Insert Table 12 here]

5.3 VAR Model with Google Trend as a Control Variable

Liu and Tsyvinski (2020) show that the investor attention significantly predicts one-week to six week ahead cumulative coin market returns. They use a weekly measure of Google search for ”Bitcoin” as a proxy for investor attention. Following Liu and Tsyvinski (2020), we construct the deviation of Google searches for the word ”Bitcoin” in a given day compared with the average of those in the preceding thirty days. We further standardize the daily deviation measure to have a mean of zero and a standard deviation of one.

Table 13 reports the results of VAR models including the lagged Google trend measures as the control variable. We confirm that the investor attention, measured in terms of Google trend measure, has a positive and significant impact on the next day bitcoin returns when it is included in the VAR model of bitcoin returns and social media sentiment. Nevertheless, compared to Table 6, results of social media sentiment remain robust, which means social media sentiment captures different and much rich information than Google searches for the word "Bitcoin."

[Insert Table 13 here]

6 Conclusion

Using a state-of-the-art NLP sentiment model and social media posts/tweets related to bitcoin from Twitter, Reddit, and Stocktwits, we investigate the relations and causality effects of social media sentiment, attention and disagreement, on bitcoin trading activity, returns, volatility and higher moments. First, we provide evidence of a reciprocal causality effect between higher social media sentiment and positive bitcoin returns, leading to a complex interplay between these two quantities. Furthermore, we showe that positive bitcoin sentiment and increased attention (proxied by the number of posts/tweets) lead to an increase in trading volume in subsequent days.

The relation between volatility and sentiment is more subtle. We do not find any evidence that sentiment directly affects volatility, although it affects daily returns kurtosis. On the other hand, we provide evidence that positive changes in social media sentiment lead to a decrease in daily realized volatility, and an increase in daily returns skewness. We further showe that higher social media attention and disagreement increase bitcoin price volatility. Overall, these findings are consistent among the three social media sources we used, although the magnitude of the impact from different social media varies.

Masked Input Text	*"Dow [mask] 900 points for worst day of year amid fears of new Covid variant."*				
Original word masked	*"fell"*				
Top 5 predicted words	"fell"	"falls"	"loses"	"rose"	"shed"
Predicted probability	0.30	0.05	0.04	0.04	0.03

Table 1: Example of the pretrained FinRoBERTa model performing the MLM task. The model is asked to predict the masked word given the rest of the input text.

Dataset \ Model	FinRoBERTa	RoBERTa Base
FPB-100	**0.9602**	0.7604
FPB-75	**0.9267**	0.7313
FPB-66	**0.8848**	0.6745

Table 2: The test accuracy of the FinRoBERTa model and the generic RoBERTa Base model on FPB dataset with different agreement levels. The lower the agreement level of the data is, the harder it is for any model to achieve high accuracy, because the sentiment in sentences with a lower agreement level are less clear even to an expert.

Text	**TextBlob**	**LM Dict**	**FinRoBERTa**
“If I use bitcoin as a store of value, transaction volume is not a very interesting metric for me. That being said, bitcoin transaction volume has been increasing when measured in terms of the goods and services that can be purchased with it.”	-0.192	0.000 (Pos: 1, Neg: 1)	0.943
“Why a Top Analyst Thinks Bitcoin Price Could Fall By 20% Before Bottoming”	0.500	0.000 (Pos: 0, Neg: 0)	-0.890
“🤷🤷 that’s why I’m here asking what the best route is to buy a bitcoin”	1.000	0.3536 (Pos: 1, Neg: 0)	-0.004
“Good entry point or y’all waiting? No Moon boys please. I’m expecting Bitcoin to correct down to 10k so OMG should drop down to $10ish as well. Thoughts?”	0.130	0.000 (Pos: 1, Neg: 1)	-0.957
”Bitcoin gets 15% down in just a day. The cryptocurrency value shows its lowest level in months. Digital currency prices fell considerably for the second consecutive day due to the impact produced by Goldman Sachs and its decision to stop its plans to launch a persistent cryptocurrency desk. Ethical hacking specialists report that the price of a unit of Bitcoin, the most widely known digital currency in the world, fell by more than $1.1k USD in a period of 24 hours, representing a decrease...”	0.061	-0.198 (Pos: 0, Neg: 2)	-0.980
“And yet Bitcoin is slowly clawing back market dominance”	-0.150	0.000 (Pos: 1, Neg: 1)	0.976
”Yes, BitMEX Liquidations Caused Bitcoin Price to Crash; Here’s How”	0.000	-0.354 (Pos: 0, Neg: 1)	-0.768

"So what's wrong with Bitcoin Cash, in terms of its technological changes? So far, all I've heard is "it's too simple" even though it, thus far, has greatly improved the usability of Bitcoin as a currency."	0.100	0.218 (Pos: 2, Neg: 1)	0.988
"Gold have like 5k years, also have industrial usage. Bitcoin is just money, nothing else. If can't be the best on that, is done."	1.000	0.267 (Pos: 1, Neg: 0)	-0.002
"They'll hit the entry points. Bitcoin is going to outpace badly in the war though. Just by being an always available alternative to an ever growing list of inflationary and manipulated currencies built under a system that heavily favors the interests of the banks and lawmakers."	-0.166	0.000 (Pos: 2, Neg: 2)	0.866

Table 3: Examples of bitcoin-related social media posts measured with 1) General sentiment by TextBlob, 2) Financial sentiment by Loughran-McDonald dictionary (LM Dict), and 3) Financial sentiment by our FinRoBERTa model. All scores range from -1 (most negative) to 1 (most positive). When using the dictionary to measure the sentiment of a text, we first delete the stopwords, i.e. extremely common words which have little value for determining the sentiment, such as "the", "he", "in", "that", etc. Then, we compute the total number of words left in the text, and count the number of positive and negative words in it according to the dictionary. Last, we calculate $p = (num\_positive\_words – num\_negative\_words)/total\_num\_words$, and the $sentiment\ score = sqr(p)$ if $p > 0$ and $-sqrt(-p)$ if $p <= 0$. In the "LM Dict" column, the $(num\_positive\_words, num\_negative\_words)$ are also shown below the $sentiment\ score$.

	Bitcoin		Sentiment (Sent)			Number (Nb)			Disagreement (Dis)		
	Ret	Vol	Twitter	Reddit	StockTwits	Twitter	Reddit	StockTwits	Twitter	Reddit	StockTwits
Min	-0.46	2.92	-0.09	-0.06	-0.04	2.43	0	0.20	0.21	0	0.23
Max	0.17	74.16	0.16	0.23	0.10	69.73	16.08	8.68	0.39	0.46	0.39
Mean	0.0007	18.62	0.04	0.02	0.03	19.45	4.10	0.99	0.28	0.26	0.29
Std	0.04	13.80	0.02	0.02	0.02	9.06	1.67	0.86	0.03	0.03	0.02
Skewness	-1.48	0.95	-0.43	1.01	-0.09	1.81	2.45	3.74	0.47	-0.39	0.26
Kurtosis	21.08	3.39	5.52	15.10	2.74	7.21	12.16	23.94	3.66	14.24	3.35

Table 4: Descriptive statistics. Volume is in millions, number of posts is in thousands.

	Returns	Twitter	Reddit	StockTwits	Traditional Media	Twitter Orth	Reddit Orth
Twitter	0.4415						
Reddit	0.3247	0.7663					
StockTwits	0.4178	0.6395	0.5130				
Traditional Media	0.1339	0.3311	0.2754	0.1756			
Twitter Orth	0.5497	0.8617	0.6704	0.5720	0.3015		
Reddit Orth	0.3751	0.6222	0.9284	0.4369	0.2404	0.7221	
StockTwits Orth	0.4555	0.5148	0.4237	0.9573	0.1410	0.5975	0.4564

Table 5: Correlation matrix for bitcoin returns, sentiments and orthogonal sentiments. The orthogonal sentiments are the residuals from the regression of sentiments on lagged bitcoin returns.

Twitter

Vars X_t, Y_t	Y_{t-1}		Y_{t-2}		X_{t-1}		X_{t-2}		$Sent^{trad}_{t-1}$	
Ret, Sent	0.2513***	(2.6693)	-0.1317**	(-1.7356)	0.1864***	(8.6359)	0.0180	(0.8355)	-0.0048	(-1.2749)
Ret, Nb	0.0034	(0.1462)	-0.0078	(-0.3342)	-0.1382***	(-2.4015)	0.0886*	(1.5410)	-0.0029	(-0.7881)
Ret, Dis	0.0609	(0.7765)	0.0869	(1.1086)	-0.0226*	(-1.3152)	-0.0218	(-1.2752)	-0.0023	(-0.6414)
Ret, Vol	0.0003	(0.0135)	0.0157	(0.7010)	-0.0250	(-0.4152)	0.0495	(0.8343)	-0.0035	(-0.9609)
Vol, Nb	0.1235***	(3.1948)	-0.1535***	(-3.9793)	-0.0344	(-0.9513)	0.0284	(0.7814)	0.0053	(0.9250)
Vol, Sent	0.1884*	(1.5578)	0.0278	(0.2392)	0.0144	(1.2822)	-0.0071	(-0.6351)	0.0031	(0.5061)
Vol, Dis	0.0464	(0.3478)	-0.2222**	(-1.6805)	0.0263***	(2.4449)	-0.0174*	(-1.6232)	0.0066	(1.1453)

Reddit

Vars X_t, Y_t	Y_{t-1}		Y_{t-2}		X_{t-1}		X_{t-2}		$Sent^{trad}_{t-1}$	
Ret, Sent	0.0838	(0.9083)	-0.0712	(-0.8553)	0.1380***	(8.2014)	0.0226*	(1.3159)	-0.0032	(-0.8511)
Ret, Nb	-0.0425*	(-1.4811)	0.0136	(0.4795)	-0.0945**	(-2.0484)	0.0092	(0.1992)	-0.0031	(-0.8630)
Ret, Dis	0.0479	(0.7494)	-0.0144	(-0.2254)	-0.0320*	(-1.5240)	-0.0239	(-1.1445)	-0.0027	(-0.7372)
Ret, Vol	0.0003	(0.0135)	0.0157	(0.7010)	-0.0250	(-0.4152)	0.0495	(0.8343)	-0.0035	(-0.9609)
Vol, Nb	0.0321	(0.6826)	-0.0829**	(-1.7720)	0.0131	(0.4623)	-0.0274	(-0.9668)	0.0064	(1.1199)
Vol, Sent	0.2032*	(1.5066)	0.0736	(0.5632)	-0.0014	(-0.1386)	0.0042	(0.4268)	0.0038	(0.6329)
Vol, Dis	0.1037	(0.9982)	-0.1464*	(-1.4207)	0.0171*	(1.3448)	-0.0175*	(-1.3704)	0.0072	(1.2587)

StockTwits

Vars X_t, Y_t	Y_{t-1}		Y_{t-2}		X_{t-1}		X_{t-2}		$Sent^{trad}_{t-1}$	
Ret, Sent	0.2302***	(2.6161)	-0.1082*	(-1.3144)	0.0813***	(4.1014)	0.0363**	(1.8552)	-0.0034	(-0.9199)
Ret, Nb	-0.0533***	(-2.5297)	0.0552***	(2.5976)	-0.0953*	(-1.4889)	0.0458	(0.7182)	-0.0024	(-0.6662)
Ret, Dis	-0.0102	(-0.1359)	-0.0781	(-1.0334)	0.0212	(1.1751)	-0.0056	(-0.3123)	-0.0035	(-0.9349)
Ret, Vol	0.0003	(0.0135)	0.0157	(0.7010)	-0.0250	(-0.4152)	0.0495	(0.8343)	-0.0035	(-0.9609)
Vol, Nb	0.1305***	(3.3520)	-0.2123***	(-5.5364)	-0.1477***	(-3.3026)	0.1424***	(3.1738)	0.0053	(0.9398)
Vol, Sent	0.1323	(1.0662)	0.0306	(0.2487)	0.0005	(0.0437)	0.0002	(0.0204)	0.0056	(0.9548)
Vol, Dis	-0.0492	(-0.4033)	0.1776*	(1.4640)	0.0030	(0.2751)	-0.0082	(-0.7522)	0.0082*	(1.4048)

Traditional Media

Vars X_t, Y_t	Y_{t-1}		Y_{t-2}		X_{t-1}		X_{t-2}	
Ret, Sent	0.0001	(0.0278)	0.0030	(0.9588)	1.3223***	(4.4484)	0.9862***	(3.2927)
Vol, Sent	0.0051	(0.9727)	0.0159***	(3.0529)	0.1405	(0.8220)	0.1053	(0.6153)

Table 6: VAR models with two lags. The two variables (e.g., "Ret, Sent") in the first column for each row represent X_t and Y_t of a VAR model defined in equations (4.1 & 4.2), correspondingly. Coefficients Y_{t-1}, Y_{t-2} are the loadings in the equation (4.1) for X_t, and coefficients X_{t-1}, X_{t-2} are the loadings in the equation (4.2) for Y_t.

Const	RV_{t-1}	Ret_{t-1}	$Sent^{T}_{t-1}$	$Sent^{R}_{t-1}$	$Sent^{S}_{t-1}$	$Sent^{trad}_{t-1}$
0.0014***	0.4934***		-0.0113**			-0.0004
(5.19)	(18.20)		(-2.13)			(-1.11)
0.0012***	0.4944***			-0.0132**		-0.0004
(6.23)	(18.28)			(-2.07)		(-1.28)
0.0015***	0.4890***				-0.0180***	-0.0004
(6.74)	(18.32)				(-3.32)	(-1.35)
0.0006**	0.4863***	-0.0284***	0.0086			-0.0004
(2.12)	(18.54)	(-8.78)	(1.54)			(-1.28)
0.0009***	0.4787***	-0.0267***		0.0022		-0.0003
(4.93)	(18.24)	(-8.66)		(0.34)		(-0.93)
0.0010***	0.4767***	-0.0264***			0.0001	-0.0003
(4.28)	(18.36)	(-8.23)			(0.02)	(-0.87)

Table 7: OLS regressions of realised variance of bitcoin intraday returns. RV is the median of daily realised variance over 9 most active exchanges. Independent variables are the lagged RV, bitcoin returns, and sentiment over the three social media sources and the traditional media. $Sent^{T}_{t-1}$, $Sent^{R}_{t-1}$, $Sent^{S}_{t-1}$, and $Sent^{trad}_{t-1}$ represent the sentiment from Twitter, Reddit, StockTwits, and the traditional media, correspondingly.

Skewness						
Const	MoM_{t-1}	Ret_{t-1}	$Sent^{T}_{t-1}$	$Sent^{R}_{t-1}$	$Sent^{S}_{t-1}$	$Sent^{trad}_{t-1}$
0.0000***	-0.2263***		-0.0008***			-0.0000
(4.51)	(-7.73)		(-4.69)			(-1.09)
0.0000***	-0.2254***			-0.0007***		-0.0000*
(2.99)	(-7.65)			(-3.37)		(-1.75)
0.0000**	-0.2172***				-0.0004**	-0.0000**
(2.26)	(-7.38)				(-2.19)	(-2.34)
0.0000	-0.1457***	-0.0011***	0.0000			-0.0000
(0.34)	(-4.98)	(-9.79)	(0.09)			(-1.28)
0.0000	-0.1456***	-0.0011***		0.0000		-0.0000
(0.54)	(-4.99)	(-10.36)		(0.12)		(-1.31)
-0.0000	-0.1375***	-0.0012***			0.0005**	-0.0000
(-1.59)	(-4.77)	(-11.02)			(2.51)	(-1.63)

Kurtosis						
Const	MoM_{t-1}	Ret_{t-1}	$Sent^{T}_{t-1}$	$Sent^{R}_{t-1}$	$Sent^{S}_{t-1}$	$Sent^{trad}_{t-1}$
0.0000***	0.2331***		-0.0001**			-0.0000
(2.79)	(7.89)		(-2.31)			(-1.33)
0.0000**	0.2372***			-0.0001		-0.0000*
(2.04)	(8.01)			(-1.31)		(-1.76)
0.0000***	0.2370***				-0.0001**	-0.0000*
(2.79)	(8.07)				(-2.29)	(-1.78)
-0.0000*	0.2472***	-0.0002***	0.0001***			-0.0000
(-1.72)	(8.87)	(-11.84)	(2.82)			(-1.58)
-0.0000	0.2448***	-0.0002***		0.0001**		-0.0000
(-0.44)	(8.78)	(-11.88)		(2.26)		(-1.29)
-0.0000	0.2415***	-0.0002***			0.0001**	-0.0000
(-1.11)	(8.73)	(-11.77)			(2.53)	(-1.08)

Table 8: OLS regressions of higher moments (skewness and kurtosis) of bitcoin intraday returns. In top panel, the dependent variable MoM is the daily realised skewness. In bottom panel, MoM is daily realised kurtosis. In both cases, a median is taken over the 9 most active bitcoin exchanges. Independent variables are the lagged values of skewness or kurtosis, bitcoin returns, and sentiment over the three social media sources and the traditional media. $Sent^{T}_{t-1}$, $Sent^{R}_{t-1}$, $Sent^{S}_{t-1}$, and $Sent^{trad}_{t-1}$ represent the sentiment from Twitter, Reddit, StockTwits, and the traditional media, correspondingly.

	Const	MoM_{t-1}	Ret_{t-1}	$Sent^T_{t-1}$	Nb^T_{t-1}	Dis^T_{t-1}
$MoM = RV_t$	-0.0037***	0.4063***	-0.0297***	0.0048	0.0751***	0.0110**
	(-2.71)	(13.53)	(-9.29)	(0.88)	(5.57)	(2.25)
$MoM = Skew_t$	-0.0001***	-0.1566***	-0.0011***	-0.0001	0.0007	0.0004***
	(-3.04)	(-5.34)	(-9.76)	(-0.50)	(1.63)	(2.85)
$MoM = Kurt_t$	-0.0000**	0.2368***	-0.0002***	0.0001**	0.0001	0.0001*
	(-2.31)	(8.38)	(-11.85)	(2.29)	(1.14)	(1.82)
	Const	MoM_{t-1}	Ret_{t-1}	$Sent^R_{t-1}$	Nb^R_{t-1}	Dis^R_{t-1}
$MoM = RV_t$	-0.0041***	0.3800***	-0.0271***	-0.0004	0.5220***	0.0120***
	(-3.51)	(13.03)	(-9.02)	(-0.06)	(6.96)	(2.65)
$MoM = Skew_t$	-0.0001***	-0.1556***	-0.0011***	-0.0002	0.0011	0.0005***
	(-3.04)	(-5.32)	(-10.15)	(-0.85)	(0.45)	(3.08)
$MoM = Kurt_t$	-0.0000**	0.2357***	-0.0002***	0.0001	0.0001	0.0001**
	(-2.12)	(8.35)	(-11.80)	(1.43)	(0.19)	(2.06)
	Const	MoM_{t-1}	Ret_{t-1}	$Sent^S_{t-1}$	Nb^S_{t-1}	Dis^S_{t-1}
$MoM = RV_t$	-0.0001	0.3838***	-0.0291***	0.0032	0.9857***	0.0005
	(-0.05)	(12.89)	(-9.02)	(0.55)	(6.08)	(0.09)
$MoM = Skew_t$	0.0001	-0.1351***	-0.0013***	0.0006***	0.0116**	-0.0003
	(1.00)	(-4.69)	(-11.46)	(3.07)	(2.47)	(-1.52)
$MoM = Kurt_t$	0.0000	0.2367***	-0.0002***	0.0001***	0.0014	-0.0001
	(1.18)	(8.51)	(-12.09)	(2.98)	(1.60)	(-1.55)

Table 9: OLS regressions of daily realised variance, skewness and kurtosis of bitcoin intraday returns. Independent variables are the lagged values of MoM, lagged bitcoin returns, lagged sentiment, lagged number of tweets/post (in millions), and lagged disagreement, taking on social media post at a time.

	Const	MoM_{t-1}	Ret_{t-1}	$\Delta Sent^{T}_{t-1}$	$\Delta Sent^{T}_{t}$	$\Delta Sent^{trad}_{t-1}$	$\Delta Sent^{trad}_{t}$
$MoM = RV_t$	0.0009***	0.4959***	-0.0258***	0.0063	-0.0269***	-0.0006**	-0.0002
	(7.59)	(19.70)	(-7.86)	(0.99)	(-4.63)	(-2.52)	(-0.98)
$MoM = Skew_t$	0.0000	-0.1507***	-0.0013***	0.0006***	0.0008***	-0.0000	0.0000
	(0.96)	(-5.29)	(-11.65)	(2.95)	(4.28)	(-0.32)	(0.55)
$MoM = Kurt_t$	0.0000*	0.2404***	-0.0002***	0.0001*	-0.0000	-0.0000	-0.0000
	(1.83)	(8.69)	(-11.19)	(1.87)	(-0.51)	(-1.61)	(-0.12)
	Const	MoM_{t-1}	Ret_{t-1}	$\Delta Sent^{R}_{t-1}$	$\Delta Sent^{R}_{t}$	$\Delta Sent^{trad}_{t-1}$	$\Delta Sent^{trad}_{t}$
$MoM = RV_t$	0.0009***	0.4875***	-0.0266***	0.0107*	-0.0168***	-0.0007***	-0.0004*
	(7.69)	(19.37)	(-8.68)	(1.67)	(-2.70)	(-2.70)	(-1.65)
$MoM = Skew_t$	0.0000	-0.1442***	-0.0012***	0.0004*	0.0003	0.0000	0.0000
	(0.92)	(-5.02)	(-11.27)	(1.71)	(1.21)	(0.24)	(1.47)
$MoM = Kurt_t$	0.0000*	0.2426***	-0.0002***	0.0001**	-0.0000	-0.0000	-0.0000
	(1.82)	(8.79)	(-11.70)	(1.96)	(-1.17)	(-1.50)	(-0.00)
	Const	MoM_{t-1}	Ret_{t-1}	$\Delta Sent^{S}_{t-1}$	$\Delta Sent^{S}_{t}$	$\Delta Sent^{trad}_{t-1}$	$\Delta Sent^{trad}_{t}$
$MoM = RV_t$	0.0009***	0.4853***	-0.0281***	0.0056	-0.0177***	-0.0007***	-0.0004*
	(7.72)	(19.26)	(-8.98)	(0.98)	(-3.25)	(-2.62)	(-1.80)
$MoM = Skew_t$	0.0000	-0.1423***	-0.0013***	0.0010***	0.0003*	0.0000	0.0000
	(0.95)	(-5.01)	(-12.25)	(4.96)	(1.91)	(0.04)	(1.37)
$MoM = Kurt_t$	0.0000*	0.2363***	-0.0002***	0.0001***	-0.0000	-0.0000	-0.0000
	(1.85)	(8.59)	(-12.30)	(3.32)	(-0.21)	(-1.56)	(-0.27)

Table 10: OLS regressions of daily realised skewness and kurtosis of bitcoin intraday returns. Independent variables are the lagged values of skewness or kurtosis, lagged bitcoin returns, and variations in sentiment defined as $\Delta Sent_t = Sent_t - Sent_{t-1}$.

Twitter

Vars X_t, Y_t	Y_{t-1}		Y_{t-2}		X_{t-1}		X_{t-2}		$Sent^{trad}_{t-1}$		Other Control Variables
Ret, Sent	0.267***	(2.85)	-0.111*	(-1.46)	0.185***	(8.59)	-0.003	(-0.14)	-0.005	(-1.21)	Y
Ret, Nb	0.006	(0.25)	-0.010	(-0.44)	-0.136***	(-2.36)	0.076	(1.24)	-0.002	(-0.63)	Y
Ret, Dis	0.054	(0.69)	0.081	(1.04)	-0.023*	(-1.35)	-0.020	(-1.10)	-0.002	(-0.49)	Y
Ret, Vol	0.001	(0.04)	0.015	(0.66)	-0.032	(-0.53)	0.035	(0.56)	-0.003	(-0.78)	Y
Vol, Nb	0.133***	(3.45)	-0.163***	(-4.24)	-0.036	(-1.00)	0.031	(0.84)	0.006	(1.01)	Y
Vol, Sent	0.175*	(1.44)	0.045	(0.38)	0.016*	(1.42)	-0.009	(-0.80)	0.004	(0.62)	Y
Vol, Dis	0.059	(0.44)	-0.265**	(-2.00)	0.025**	(2.33)	-0.016*	(-1.51)	0.007	(1.24)	Y

Reddit

Vars X_t, Y_t	Y_{t-1}		Y_{t-2}		X_{t-1}		X_{t-2}		$Sent^{trad}_{t-1}$		Other Control Variables
Ret, Sent	0.101	(1.11)	-0.052	(-0.62)	0.137***	(8.17)	0.007	(0.39)	-0.003	(-0.76)	Y
Ret, Nb	-0.042*	(-1.48)	0.014	(0.48)	-0.086**	(-1.87)	-0.002	(-0.04)	-0.003	(-0.70)	Y
Ret, Dis	0.042	(0.65)	-0.006	(-0.10)	-0.032*	(-1.53)	-0.026	(-1.16)	-0.002	(-0.57)	Y
Ret, Vol	0.001	(0.04)	0.015	(0.66)	-0.032	(-0.53)	0.035	(0.56)	-0.003	(-0.78)	Y
Vol, Nb	0.033	(0.71)	-0.086**	(-1.84)	0.019	(0.66)	-0.032	(-1.12)	0.007	(1.21)	Y
Vol, Sent	0.203*	(1.50)	0.095	(0.72)	-0.000	(-0.03)	0.003	(0.31)	0.004	(0.71)	Y
Vol, Dis	0.112	(1.07)	-0.168*	(-1.63)	0.017*	(1.34)	-0.018*	(-1.37)	0.008*	(1.35)	Y

StockTwits

Vars X_t, Y_t	Y_{t-1}		Y_{t-2}		X_{t-1}		X_{t-2}		$Sent^{trad}_{t-1}$		Other Control Variables
Ret, Sent	0.244***	(2.79)	-0.068	(-0.82)	0.083***	(4.23)	0.014	(0.70)	-0.003	(-0.79)	Y
Ret, Nb	-0.052***	(-2.48)	0.054***	(2.56)	-0.089*	(-1.38)	0.030	(0.43)	-0.002	(-0.51)	Y
Ret, Dis	-0.006	(-0.08)	-0.095	(-1.25)	0.018	(1.03)	0.012	(0.61)	-0.003	(-0.78)	Y
Ret, Vol	0.001	(0.04)	0.015	(0.66)	-0.032	(-0.53)	0.035	(0.56)	-0.003	(-0.78)	Y
Vol, Nb	0.137***	(3.53)	-0.220***	(-5.76)	-0.146***	(-3.24)	0.141***	(3.12)	0.006	(1.05)	Y
Vol, Sent	0.127	(1.02)	0.036	(0.28)	0.003	(0.27)	-0.002	(-0.21)	0.006	(1.05)	Y
Vol, Dis	-0.067	(-0.55)	0.163*	(1.33)	0.002	(0.21)	-0.008	(-0.71)	0.008*	(1.44)	Y

Table 11: VAR models with two lags with financial control variables. The two variables (e.g., "Ret, Sent") in the first column for each row represent X_t and Y_t of a VAR model defined in equations (5.1 & 5.2), correspondingly. Coefficients Y_{t-1}, Y_{t-2} are the loadings in the equation (5.1) for X_t, and coefficients X_{t-1}, X_{t-2} are the loadings in the equation (5.2) for Y_t.

	Const	MoM_{t-1}	Ret_{t-1}	$Sent^{C1}_{t-1}$	$Sent^{C2}_{t-1}$	$Sent^{C3}_{t-1}$
$MoM = RV_t$	0.0010***	0.4831***	-0.0279***	0.0026	-0.0047	-0.0176
	(7.67)	(18.29)	(-8.45)	(0.65)	(-0.63)	(-1.58)
$MoM = Skew_t$	0.0000	-0.1445***	-0.0012***	0.0001	0.0007***	0.0003
	(0.92)	(-4.95)	(-10.19)	(0.70)	(2.77)	(0.94)
$MoM = Kurt_t$	0.0000*	0.2481***	-0.0002***	0.0001***	0.0000	-0.0000
	(1.82)	(8.88)	(-11.81)	(2.75)	(0.48)	(-0.16)
	Const	MoM_{t-1}	Ret_{t-1}	Dis^{C1}_{t-1}	Dis^{C2}_{t-1}	Dis^{C3}_{t-1}
$MoM = RV_t$	0.0010***	0.4483***	-0.0283***	0.0099***	-0.0111**	-0.0002
	(8.16)	(16.44)	(-9.49)	(2.71)	(-2.13)	(-0.02)
$MoM = Skew_t$	0.0000	-0.1515***	-0.0011***	0.0004***	-0.0004**	0.0000
	(0.92)	(-5.29)	(-11.52)	(3.21)	(-2.26)	(0.02)
$MoM = Kurt_t$	0.0000*	0.2233***	-0.0002***	0.0001***	-0.0001**	-0.0000
	(1.85)	(8.01)	(-12.18)	(2.59)	(-1.97)	(-0.23)
	Const	MoM_{t-1}	Ret_{t-1}	Nb^{C1}_{t-1}	Nb^{C2}_{t-1}	Nb^{C3}_{t-1}
$MoM = RV_t$	0.0012***	0.3626***	-0.0278***	0.0863***	0.4900***	0.4307**
	(9.46)	(12.31)	(-9.63)	(6.42)	(4.80)	(2.31)
$MoM = Skew_t$	0.0000	-0.1478***	-0.0011***	0.0009**	-0.0023	0.0123**
	(0.92)	(-5.16)	(-11.33)	(2.08)	(-0.70)	(2.03)
$MoM = Kurt_t$	0.0000*	0.2291***	-0.0002***	0.0001	-0.0005	0.0019*
	(1.84)	(8.21)	(-12.03)	(1.10)	(-0.82)	(1.70)

Table 12: OLS regressions of daily realised variance, skewness and kurtosis of bitcoin intraday returns. Independent variables are the lagged values of MoM, lagged bitcoin returns, lagged first three principal components of sentiment (C1 to C3).

Twitter

Vars X_t, Y_t	Y_{t-1}		Y_{t-2}		X_{t-1}		X_{t-2}		$Sent^{trad}_{t-1}$		$Goog_{t-1}$	
Ret, Sent	0.255***	(3.06)	-0.103*	(-1.58)	0.198***	(11.77)	-0.002	(-0.11)	-0.002	(-0.77)	0.005***	(2.41)
Ret, Nb	-0.007	(-0.42)	0.012	(0.70)	-0.132***	(-2.59)	0.051	(1.00)	0.000	(0.01)	0.005***	(2.65)
Ret, Dis	0.029	(0.50)	0.009	(0.15)	-0.038**	(-2.33)	-0.000	(-0.00)	-0.000	(-0.02)	0.004**	(2.00)
Ret, Vol	-0.013	(-0.73)	0.017	(0.97)	0.025	(0.49)	0.061	(1.18)	-0.000	(-0.11)	0.005***	(2.36)
Vol, Nb	0.113***	(3.62)	-0.164***	(-5.24)	0.042*	(1.34)	-0.022	(-0.69)	0.006	(1.08)	-0.010***	(-2.83)
Vol, Sent	0.198**	(1.80)	-0.005	(-0.04)	0.008	(0.92)	-0.005	(-0.58)	0.004	(0.65)	-0.007**	(-1.85)
Vol, Dis	-0.035	(-0.34)	-0.220**	(-2.19)	0.037***	(3.72)	-0.025***	(-2.45)	0.007*	(1.31)	-0.004	(-1.00)

Reddit

Vars X_t, Y_t	Y_{t-1}		Y_{t-2}		X_{t-1}		X_{t-2}		$Sent^{trad}_{t-1}$		$Goog_{t-1}$	
Ret, Sent	0.121*	(1.58)	-0.011	(-0.16)	0.135***	(9.84)	0.016	(1.12)	-0.001	(-0.35)	0.005***	(2.35)
Ret, Nb	-0.038**	(-1.73)	0.025	(1.12)	-0.107***	(-2.63)	0.025	(0.62)	0.000	(0.01)	0.005***	(2.47)
Ret, Dis	0.002	(0.03)	-0.015	(-0.31)	-0.044***	(-2.40)	-0.021	(-1.14)	-0.000	(-0.04)	0.005***	(2.44)
Ret, Vol	-0.013	(-0.73)	0.017	(0.97)	0.025	(0.49)	0.061	(1.18)	-0.000	(-0.11)	0.005***	(2.36)
Vol, Nb	0.084**	(2.23)	-0.163***	(-4.24)	0.059***	(2.39)	-0.060***	(-2.44)	0.005	(1.04)	-0.011***	(-2.88)
Vol, Sent	0.283***	(2.39)	-0.047	(-0.41)	0.004	(0.46)	-0.003	(-0.37)	0.004	(0.73)	-0.007**	(-1.98)
Vol, Dis	0.037	(0.43)	-0.171**	(-2.01)	0.031***	(2.86)	-0.039***	(-3.57)	0.007*	(1.38)	-0.005*	(-1.30)

StockTwits

Vars X_t, Y_t	Y_{t-1}		Y_{t-2}		X_{t-1}		X_{t-2}		$Sent^{trad}_{t-1}$		$Goog_{t-1}$	
Ret, Sent	0.200***	(2.95)	-0.080	(-1.25)	0.079***	(4.75)	0.007	(0.42)	-0.001	(-0.22)	0.005***	(2.65)
Ret, Nb	-0.037**	(-2.18)	0.036**	(2.17)	-0.098**	(-1.81)	0.081*	(1.49)	0.000	(0.03)	0.006***	(2.67)
Ret, Dis	0.011	(0.20)	-0.041	(-0.73)	0.008	(0.46)	0.008	(0.47)	-0.000	(-0.09)	0.005***	(2.52)
Ret, Vol	-0.013	(-0.73)	0.017	(0.97)	0.025	(0.49)	0.061	(1.18)	-0.000	(-0.11)	0.005***	(2.36)
Vol, Nb	0.152***	(4.56)	-0.241***	(-7.50)	-0.150***	(-4.06)	0.169***	(4.56)	0.006	(1.11)	-0.004	(-1.12)
Vol, Sent	0.067	(0.65)	0.011	(0.11)	0.005	(0.52)	-0.007	(-0.84)	0.007	(1.25)	-0.007**	(-1.85)
Vol, Dis	-0.168**	(-1.78)	0.157**	(1.66)	0.013*	(1.35)	-0.024***	(-2.57)	0.007*	(1.43)	-0.006**	(-1.80)

Table 13: VAR models with two lags. The two variables (e.g., "Ret, Sent") in the first column for each row represent X_t and Y_t of a VAR model defined in equations (5.1 & 5.2), correspondingly. Coefficients Y_{t-1}, Y_{t-2} are the loadings in the equation (5.1) for X_t, and coefficients X_{t-1}, X_{t-2} are the loadings in the equation (5.2) for Y_t. $Goog_{t-1}$ is the lagged Google trend measure.

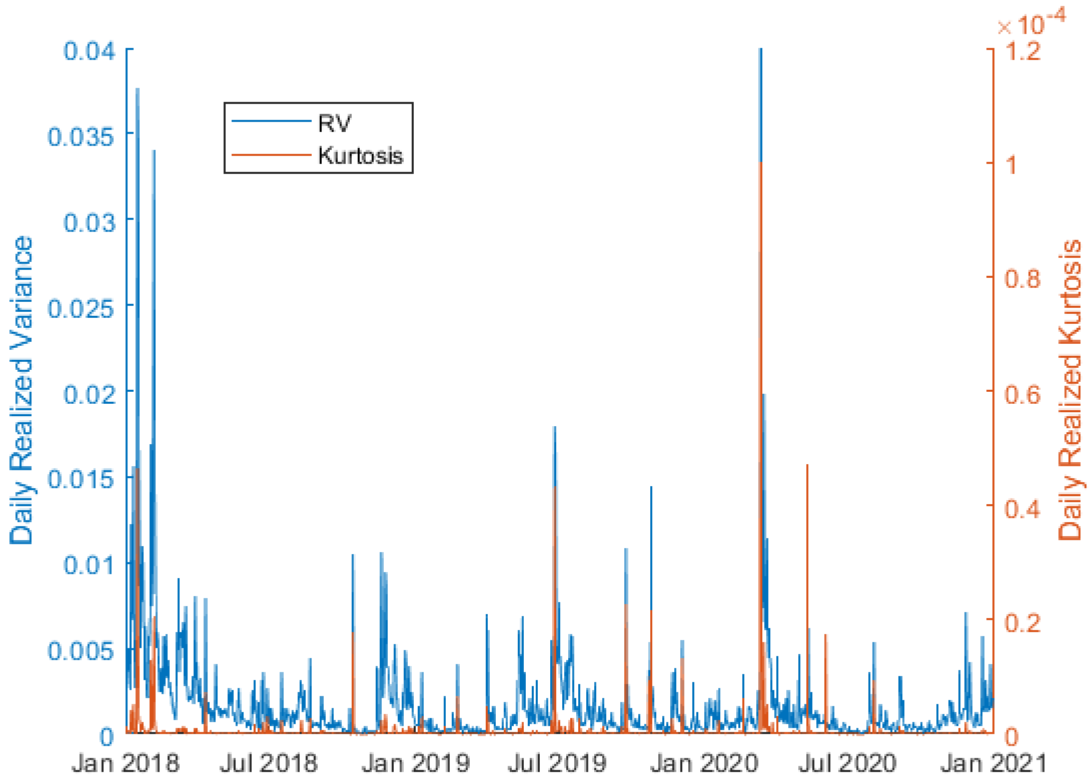

Figure 1: This figure plots daily realized variance and kurtosis, which are estimated from intraday 5-minute bitcoin returns.

References

Agrawal, S., Azar, P.D., Lo, A.W., Singh, T., 2018. Momentum, mean-reversion, and social media: Evidence from stocktwits and twitter. The Journal of Portfolio Management 44, 85–95.

Allee, K.D., Deangelis, M.D., 2015. The structure of voluntary disclosure narratives: Evidence from tone dispersion. Journal of Accounting Research 53, 241–274.

Antweiler, W., Frank, M.Z., 2004. Is all that talk just noise? the information content of internet stock message boards. The Journal of finance 59, 1259–1294.

Araci, D., 2019. FinBERT: Financial sentiment analysis with pre-trained language models. arXiv preprint arXiv:1908.10063 .

Awais, M., Yang, J., 2021. Does divergence of opinions make better minds? evidence from social media. Working paper .

Back, M.D., Stopfer, J.M., Vazire, S., Gaddis, S., Schmukle, S.C., Egloff, B., Gosling, S.D., 2010. Facebook profiles reflect actual personality, not self-idealization. Psychological Science 21, 372–374.

Bahdanau, D., Cho, K., Bengio, Y., 2014. Neural machine translation by jointly learning to align and translate.

Bartov, E., Faurel, L., Mohanram, P.S., 2017. Can Twitter Help Predict Firm-Level Earnings and Stock Returns? The Accounting Review 93, 25–57.

Betzer, A., Harries, J.P., 2021. How online comments affect stock trading - the case of gamestop. Financial Markets and Portfolio Management, forthcoming .

Blankespoor, E., Miller, G.S., White, H.D., 2014. The role of dissemination in market liquidity: Evidence from firms' use of twitter™. The Accounting Review 89, 79–112.

Bollerslev, T., Litvinova, J., Tauchen, G., 2006. Leverage and Volatility Feedback Effects in High-Frequency Data. Journal of Financial Econometrics 4, 353–384.

Brown, T.B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., 2020. Language models are few-shot learners. arXiv preprint arXiv:2005.14165 .

Cade, N.L., 2018. Corporate social media: How two-way disclosure channels influence investors. Accounting, Organizations and Society 68-69, 63–79.

Carr, P., Wu, L., 2017. Leverage effect, volatility feedback, and self-exciting market disruptions. Journal of Financial and Quantitative Analysis 52, 2119–2156.

Chen, H., De, P., Hu, Y., Hwang, B.H., 2014. Wisdom of Crowds: The Value of Stock Opinions Transmitted Through Social Media. The Review of Financial Studies 27, 1367–1403.

Cho, K., van Merrienboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., Bengio, Y., 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation, in: Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), Association for Computational Linguistics. pp. 1724–1734.

Da, Z., Engelberg, J., Gao, P., 2011. In search of attention. The Journal of Finance 66, 1461–1499.

Das, S.R., Chen, M.Y., 2007. Yahoo! for amazon: Sentiment extraction from small talk on the web. Management Science 53, 1375–1388.

DeMarzo, P.M., Vayanos, D., Zwiebel, J., 2003. Persuasion bias, social influence, and unidimensional opinions. The Quarterly Journal of Economics 118, 909–968.

Devlin, J., Chang, M.W., Lee, K., Toutanova, K., 2019. BERT: Pre-training of deep bidirectional transformers for language understanding, in: NAACL, Association for Computational Linguistics. p. 4171–4186.

Diangson, B., Jung, N., 2021. Bet if on reddit: The effects of reddit chatter on highly shorted stocks. Working paper .

Engelberg, J., Reed, A.V., Ringgenberg, M.C., 2012. How are shorts informed?: Short sellers, news, and information processing. Journal of Financial Economics 105, 260–278.

Garcia, D., 2013. Sentiment during recessions. The Journal of Finance 68, 1267–1300.

Hirshleifer, D., 2020. Presidential address: Social transmission bias in economics and finance. The Journal of Finance 75, 1779–1831.

Hobson, J.L., Mayew, W.J., Venkatachalam, M., 2012. Analyzing speech to detect financial misreporting. Journal of Accounting Research 50, 349–392.

Hu, D., Jones, C.M., Zhang, V., Zhang, X., 2021. The rise of reddit: How social media affects retail investors and short-sellers' roles in price discovery. Working paper .

Huang, A.H., Zang, A.Y., Zheng, R., 2014. Evidence on the information content of text in analyst reports. The Accounting Review 89, 2151–2180.

Huang, K., Altosaar, J., Ranganath, R., 2019. Clinicalbert: Modeling clinical notes and predicting hospital readmission. arXiv preprint arXiv:1904.05342 .

Keynes, J.M., 1936. The general theory of employment, interest and money. London: Macmillan .

Kothari, S.P., Li, X., Short, J.E., 2009. The effect of disclosures by management, analysts, and business press on cost of capital, return volatility, and analyst forecasts: A study using content analysis. The Accounting Review 84, 1639–1670.

Larcker, D.F., Zakolyukina, A.A., 2012. Detecting deceptive discussions in conference calls. Journal of Accounting Research 50, 495–540.

Lee, J., Yoon, W., Kim, S., Kim, D., Kim, S., So, C.H., Kang, J., 2019. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. Bioinformatics 36, 1234–1240.

Lee, L.F., Hutton, A.P., Shu, S., 2015. The role of social media in the capital market: Evidence from consumer product recalls. Journal of Accounting Research 53, 367–404.

Liu, F., Weng, F., Jiang, X., 2012. A broad-coverage normalization system for social media language, in: Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 1035–1044.

Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V., 2019. RoBERTa: A robustly optimized BERT pretraining approach. arXiv preprint arXiv:1907.11692 .

Liu, Y., Tsyvinski, A., 2020. Risks and Returns of Cryptocurrency. The Review of Financial Studies 34, 2689–2727.

Loughran, T., McDonald, B., 2016. Textual analysis in accounting and finance: A survey. Journal of Accounting Research 54, 1187–1230.

Loughran, T.I.M., McDonald, B., 2011. When is a liability not a liability? textual analysis, dictionaries, and 10-Ks. The Journal of Finance 66, 35–65.

Malo, P., Sinha, A., Korhonen, P., Wallenius, J., Takala, P., 2014. Good debt or bad debt: Detecting semantic orientations in economic texts. Journal of the Association for Information Science and Technology 65, 782–796.

McVay, S.E., 2006. Earnings management using classification shifting: An examination of core earnings and special items. The Accounting Review 81, 501–531.

Mikolov, T., Chen, K., Corrado, G., Dean, J., 2013. Efficient estimation of word representations in vector space.

Park, G.and Schwartz, H.A., Eichstaedt, J. C.and Kern, M.L., Kosinski, M., Stillwell, D. J.and Ungar, L.H., Seligman, M.E., 2015. Automatic personality assessment through social media language. Journal of personality and social psychology 108, 934–952.

Pedersen, L.H., 2021. Game on: Social networks and markets. Working Paper .

Pennington, J., Socher, R., Manning, C.D., 2014. Glove: Global vectors for word representation, in: Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP), pp. 1532–1543.

Radford, A., Narasimhan, K., Salimans, T., Sutskever, I., 2018. Improving language understanding by generative pre-training.

Schrand, C.M., Walther, B.R., 2000. Strategic benchmarks in earnings announcements: The selective disclosure of prior-period earnings components. The Accounting Review 75, 151–177.

Tang, V.W., 2018. Wisdom of crowds: Cross-sectional variation in the informativeness of third-party-generated product information on twitter. Journal of Accounting Research 56, 989–1034.

Tetlock, P.C., 2007. Giving content to investor sentiment: The role of media in the stock market. The Journal of Finance 62, 1139–1168.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I., 2017. Attention is all you need, in: Proceeding of the Conference on Neural Information Processing System (NIPS 2017), pp. 5998–6008.

Wolf, T., Chaumond, J., Debut, L., Sanh, V., Delangue, C., Moi, A., Cistac, P., Funtowicz, M., Davison, J., Shleifer, S., . Transformers: State-of-the-art natural language processing, in: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, pp. 38–45.

Yang, Y., Uy, M.C.S., Huang, A., 2020. FinBERT: A pretrained language model for financial communications. arXiv preprint arXiv:2006.08097 .