

Jan 17th, 12:00 AM

The Predictive Power of Social Media Sentiment for Short-Term Stock Movements

Moritz Victor Wilksch
Universität Potsdam, Germany, wilksch@uni-potsdam.de

Olga Dr. Abramova
Universität Potsdam, Germany, olga.abramova@uni-potsdam.de

Follow this and additional works at: <https://aisel.aisnet.org/wi2022>

Recommended Citation

Wilksch, Moritz Victor and Dr. Abramova, Olga, "The Predictive Power of Social Media Sentiment for Short-Term Stock Movements" (2022). *Wirtschaftsinformatik 2022 Proceedings*. 38.
https://aisel.aisnet.org/wi2022/student_track/student_track/38

This material is brought to you by the Wirtschaftsinformatik at AIS Electronic Library (AISeL). It has been accepted for inclusion in Wirtschaftsinformatik 2022 Proceedings by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact elibrary@aisnet.org.

The Predictive Power of Social Media Sentiment for Short-Term Stock Movements

Moritz Wilksch¹, Dr. Olga Abramova¹

¹University of Potsdam, Business Informatics, Potsdam, Germany
{wilksch, oabramov}@uni-potsdam.de

Abstract. Ever since modern-day financial markets existed, people have been trying to forecast movements in stock prices, as accurate predictions would entail economic benefits and the reduction of risks. This paper examines whether social media sentiment can be used to predict short-term stock movements. Using more than two years of data from Twitter, we assess the effect the extracted sentiment holds for 10 companies listed in the S&P500. Applying different sentiment analysis approaches and forecasting models, we find that for three out of the ten companies, sentiment does significantly improve the forecasting performance. A custom-built sentiment model outperforms an off-the-shelf VADER model, and tree-based models deliver better performance than linear ones. On the theoretical front, this provides evidence against the Efficient Market Hypothesis and warrants future research regarding the circumstances under which stock returns might be predictable.

Keywords: Social sentiment, Twitter, stock market, predictive power, forecasting

1 Introduction

Forecasting future returns of stocks has been an active area of research ever since the advent of modern financial markets. Most practitioners approach this problem in one of two ways: By analyzing fundamental data like balance sheets and cash flow statements to find undervalued companies to buy for a long-term investment or by conducting technical analysis to discover and exploit short-term trends and patterns in historical price movements. According to the efficient market hypothesis (EMH), neither of these approaches works to outperform the general market. The EMH, in its strong form, states that “security prices at any point in time fully reflect all available information” [1, p. 388]. This implies that generating excess returns can only be done by taking on extra risks, and thus no technique or forecasting model could elevate returns above the level of the general market at the same level of risk. In recent years, however, the EMH has been challenged for the far-from-reality assumptions and inability to explain certain phenomena. For example, the observed volatility in equity markets is higher than what would be expected under an efficient market model [2].

An alternative theory of Behavioral Finance (BF) acknowledges that market participants are subject to a wide range of cognitive biases like overconfidence and the use of heuristics [3]. Thus, BF allows for market inefficiencies caused by human behavior, e.g., overreaction or underreaction to the news. This is especially relevant since the number of retail investors and traders is rising [4] together with the expansion of online investing communities. The latter often emerge on the existing social networking sites (e.g., Reddit or Twitter), becoming a place of active exchange of opinions. As such, in January 2021, a group of individual investors organized in a community on the social media platform Reddit even managed to cause a short squeeze in GameStop stock by collectively driving its price up [5]. Further, Elon Musk, CEO of Tesla Inc., caused a sudden, 16% price jump in Bitcoin after tweeting Tesla would accept the cryptocurrency as payment for vehicles [6]. Past Information Systems (IS) studies mainly focus on the link between sentiment and information dissemination and register that emotionally charged content is shared more often in the news [7][8] and political domain [9]. Overall positive sentiment increased information sharing during a crisis event [10]. Contrary to them, in online health communities (OHC), more negative content receives greater support [11]. Considering the examples above and the fact that millions of investment-related discussions by ordinary users are available daily, a question arises:

RQ: *Does the crowd's opinion on SM hold significant predictive power regarding short-term stock movements? And if so, to what extent?*

To answer it, we study ten large-cap US companies, which are a matter of regular financial discussions on Twitter. We build our models based on the daily data scraped over a 28-month period. In contrast to most previous research in the field, e.g., [22-23], we use and compare two different sentiment analysis models: VADER, a lexicon-based model, and a custom-built sentiment analysis classifier built on a corpus of 3,000 manually labeled tweets. Moreover, we do not only use linear VARX models, but also non-linear tree-based models to predict the following days' binarized return (positive or negative) from daily sentiment scores and financial indicators. With the rigorous approach towards sentiment analysis and modeling, this work aims to contribute to the body of literature concerning BF and the weak and semi-strong form of EMH as well as to highlight the difference between the two sentiment analysis approaches. Practitioners might use findings of the modeling process to further improve existing models by including sentiment-based features and rule out model configurations with sub-par performance.

2 Theoretical Background and Conceptual Framework

2.1 Financial Markets Theories and Information Arrivals

Ever since being proposed by Eugene Fama in 1965, EMH has been a predominant model in financial theory [1]. The EMH is concerned with whether prices of securities at any given point in time fully reflect a particular subset of information. According to Fama [1], a market is efficient per definition if (1) there are no transaction costs, (2) all

information is available to all market participants at no cost, and (3) all market participants agree on the implication of current information. As these assumptions seem hardly feasible in real settings, Fama points out that these are only *sufficient* conditions; that is, they are not necessary and weaker forms of efficient markets exist. Specifically, he conceptualizes three forms of market efficiency: weak form, semi-strong form, and strong form.

In weak-form market efficiency, the subset of information that prices “fully reflect” is historical price information. Assuming that successive returns are identically distributed, this hypothesis can be expressed as

$$f(r_{j,t+1}|\Phi_t) = f(r_{j,t+1}) \quad (1)$$

where $r_{j,t+1}$ denotes the return of security j at time $t+1$, Φ_t denotes the available information at time t , and f is the probability density of the return distribution. This model is also called the *random walk* model, as the conditional independence stated in (1) implies that security prices follow a random walk. Under these assumptions, no historical price information can be used to forecast future stock returns. This also implies that the practice of technical analysis – the study of chart patterns – cannot be used to generate excess returns. In its semi-strong form, the EMH assumes that prices reflect all *publicly available* information. While this assumption is harder to test than the weak-form EMH, [12] have conducted a series of tests examining market reactions to news like stock splits, dividend- and earnings announcements. They find evidence that markets react immediately and efficiently to such events. Finally, the strong form of the EMH assumes that prices fully reflect *all information*, implying no individual can expect higher profits than the competition due to monopolistic access to information. However, there is evidence that this assumption is unrealistic as insider trading does indeed occur and generate excess returns [13].

In contrast to the EMH, Behavioral Finance (BF) does not use the simplified assumption of rational agents to describe financial markets, market participants, and their interactions with one another. Rather it is the study of how psychology impacts the decision-making of individuals. Foundations for this area of research were laid in the 1970s when psychologists Daniel Kahneman and Amos Tversky started studying judgment under uncertainty and found that humans employ certain heuristics and mental operations when assessing uncertain situations, which lead to systematic and predictable errors [14]. Subsequent research shows the effect of emotions and group behavior on decision-making processes. According to [15], there are three classes of findings in the behavioral finance literature. First, there is a catalog of biases that human decision-makers are subject to. Second, there are speculative dynamics in asset prices, where “systematic errors of unsophisticated investors [...] create profit opportunities for experts” [15, p. 9]. This also implies that the opinions and sentiment of such investors could be used by experts to gauge the emergence of price bubbles. Finally, there are findings regarding how decision processes influence decision outcomes. This is especially applicable to corporate settings, where formal decision processes are codified. Overall, unlike in neoclassical finance theory, BF does not have a unified theoretical core [15]. Instead, it is a collection of psychological models applied to

economics and attempts to explain empirical market phenomena through the behavior of individuals [16].

2.2 Empirical Evidence on the Predictive Power of Sentiment

While sentiment expressed in publicly released news articles exhibited effects on stock returns (e.g., [17]; [18]), recently, the opinion of the crowd gained importance. As such, [19] evidenced that social media sentiment has a stronger relationship with stock returns than sentiment extracted from traditional media. The most common starting point for social sentiment analysis is big social media platforms, as anyone can just sign up and start posting to a vast audience. Consequently, studies aimed to find out whether public sentiment does indeed hold such predictive power mushroomed. A happiness index calculated from Facebook posts has been shown to predict daily returns and trading volume [20-21]. Sentiment extracted from the micro-blogging platform Twitter has also been successfully used to predict short-term stock returns [22-23]. Even the opinions of users in online investing forums have been shown to predict future closing prices [24] or improve predictive power if combined with other sources [25]. The forum “Stock Twits” seems to be of particular interest, as users can label their posts as either bullish or bearish, thus providing researchers with an abundance of labeled data. In all other cases, *sentiment analysis* is the preferred technique applied on social media posts. Sentiment analysis techniques employed in forecasting stock returns most often rely on a lexicon that assigns sentiment scores to single words and aggregates them [22-23] [26]. Occasionally, researchers employ self-trained machine learning models [27], but most previous works used pre-trained models [28].

Overall, there is evidence that social sentiment can be used to forecast the future return of some securities. Sentiment extracted from Twitter using the Profile of Mood States lexicon has been used as a feature in neural networks to achieve more than 75% directional accuracy for daily forecasts of Dow Jones values [22-23]. Moreover, Twitter sentiment has been shown to granger-cause stock market return [29], which can be used by machine learning models to predict future returns of the Dow Jones and NASDAQ indexes with high accuracy [30]. [31] show that even linear models can exploit Twitter sentiment to explain a significant amount of variance in the daily returns of 69 different technology companies. Similar results are reported by [32], who not only achieve above 80% prediction accuracy but also show that adding sentiment features improves accuracy by 18%p for the Chinese SSE50 index. Even when not added to an existing financial model, sentiment indicators on their own can hold predictive power, as [33] illustrate for the Chinese stock market.

However, the results are not entirely consistent: When applied to the Bitcoin market and the corresponding online forum bitcointalk.org, [26] finds that sentiment is mainly determined through past performance and only carries limited information for price forecasting. This is confirmed by an analysis conducted by [34]. It suggests that the effect returns have on sentiment is much larger than vice versa, and although predictive power can be found for a small percentage of stocks, there is no clear pattern under which circumstances this is the case. Similarly, [27] states that stock returns and sentiment of five US technology stocks are highly correlated, but he was unable to use

this for prediction purposes. Sometimes, even if statistically significant predictive power can be found, it lacks practical significance: Acting upon predictions entails brokerage fees which often make trading strategies with minuscule upside unprofitable [35]. While social sentiment has been shown to improve stock *volatility* forecasts [28][35] and predict future trading *volume* [36], there is not yet a consensus as to whether or under which specific circumstances social sentiment holds predictive power regarding stock returns.

2.3 Conceptual Framework

Drawing on previous research on sentiment analysis of social media content and its relationship to stock returns [22-23], this paper examines the predictive power social media sentiment holds for future stock returns. To capture this emotional factor as detailed as possible, we will use not only a univariate sentiment score but also a measure of how polarized the sentiment is on any given day. Should public sentiment hold any predictive power, this would provide evidence against the weak form of the EMH as presented in equation (1) and indicate that BF might be better suited to explain modern financial markets. To make results more comparable to other research modeling future stock returns, we add basic financial indicators that are typically used to predict returns [31]. This results in the conceptual framework presented in Figure 1.




Figure 1: Conceptual framework

3 Materials and Methods

To operationalize public sentiment, we rely on the data from Twitter, one of the largest micro-blogging platforms with over 330 million monthly active users [37]. Here, the so-called cashtags, i.e., tags consisting of a “\$” sign followed by a stock ticker, conveniently reference publicly listed companies when talking about them in a financial context. This characteristic makes cashtags a working filtering mechanism to find tweets discussing investments in specific companies. Figure 2 exhibits the study flow. We began with a pre-study, scraping tweets containing for all S&P 500 companies *twint* [38]. Ultimately, we decided to choose 10 most widely discussed companies on Twitter, adjusting for shutdowns due to COVID-19. The final sample consists of large-cap technology and semiconductor companies: Tesla Inc., Apple Inc., Amazon.com Inc., Facebook Inc., Microsoft Corporation, Twitter Inc., Advanced Micro Devices Inc., Netflix Inc., Nvidia Corporation, and Intel Corporation.




Figure 2: Study procedure

In the next stage, we scraped Twitter and financial data for our sample, followed by the necessary cleaning and transformation procedures. Sentiment analysis was performed by using 1) a dictionary-based VADER model and 2) a custom sentiment model. On the modeling stage, we predicted the next day's return with: 1) VARX, 2) Random Forest, and 3) LightGBM models. Due to space limitations, we extensively disclose the methodological procedure for an interested reader under a link in the Open Science Framework (OSF) repository¹, precisely: Pre-study (Appendix A), Data Collection and Data Preprocessing (Appendix B), Sentiment Analysis Procedure (Appendix C), Feature Engineering and Data Set Characteristics (Appendix D), Forecasting Models (Appendix E), Model Evaluation (Appendix F) as well as

¹ https://osf.io/3tyaj/?view_only=c7ef10a089fb405caf58abe4cd2fae8b

intermediate results of Sentiment Analysis (Appendix G), VARX Modeling (Appendix H) and Non-linear Models, i.e., Random Forest and LightGBM (Appendix I).

4 Results

To assess how much predictive power can be attributed to the sentiment-based variables, we contrast the model of direction of daily returns (dependent variable) explained by sentiment and financial features (independent variables) vs. the model of daily returns rise or fall explained by financial features only. Each company is analyzed separately, and the final results encompass the best model using sentiment features (i.e., VADER or custom-built sentiment analysis classifier), for details on checking assumptions and lag selection, see Appendix H.

4.1 VARX Modeling

Only for five (AMD, NFLX, INTC, TSLA, MSFT) out of ten companies can a model that outperforms the baseline be found (Figure 3 and Appendix H). For these five companies, the models using VADER sentiment seem to perform slightly better than the models using the custom sentiment approach: Three out of the five working models use VADER sentiment (INTC, TSLA, MSFT), two (AMD, NFLX) use the custom sentiment scores (Figure 3). Remarkable is the case of Intel: Using VADER sentiment, the model was able to achieve an accuracy improvement of more than 11%p. However, besides Intel, the overall performance improvement is small in magnitude (3-5%p). The VARX models for Amazon, Nvidia, Facebook, and Twitter perform worse than the corresponding baselines by a large margin, indicating that even linear VARX models can overfit the training data and not generalize well.




Figure 3: Accuracy improvement over baseline, colored by sentiment method

4.2 Non-linear Models

Among the non-linear models, we tried using Random Forests and LightGBM. To assess how much predictive power can be attributed to the sentiment-based variables, we repeat the modeling process, this time removing all sentiment-related features (Appendix I).

Sentiment:	Random Forest		LightGBM		Baseline
	ml_sentiment	VADER	ml_sentiment	VADER	
TSLA	0.478	0.478	0.444	0.556*	0.494
AAPL	0.511	0.511	0.511	0.511	0.529
AMZN	0.511	0.50	0.467	0.511	0.552
FB	0.522	0.522	0.467	0.433	0.54
MSFT	0.489	0.489	0.60*	0.50	0.506
TWTR	0.533	0.533	0.533	0.544	0.552
AMD	0.467	0.411	0.544*	0.511	0.483
NFLX	0.644*	0.556	0.633	0.589	0.575
NVDA	0.589	0.589	0.589	0.60	0.609
INTC	0.489	0.589	0.60*	0.489	0.506

Table 1: Test set accuracy. Bold = beats baseline, * = best value for ticker

As the absolute accuracy scores (Table 1) can be misleading in the case of non-uniform prior class distribution, Figure 4 visualizes these scores relative to the company’s baseline accuracy (sorted by “with sentiment” model performance improvement from left to right). For six out of the ten companies, models with significant predictive power can be found (Table 1), although, for Nvidia, this only holds for the model without any sentiment features (Appendix I). For the other four companies, comparing model performance is unnecessary, as the models have not learned any generalizable pattern that beats predicting the majority class and are thus considered inadequate. It occurs that for Intel, the sentiment data does not seem to improve predictive power as the model without any sentiment data is more accurate. The same holds for AMD and Nvidia: while only using financial features yields a model with an accuracy of around 55.6% and 66.7% respectively, adding sentiment feature deteriorates the models’ performance. In contrast to this, for Microsoft, Netflix, and Tesla, the models with sentiment features achieve significantly higher performance than their counterparts. In these cases, the sentiment data hold predictive power that can be exploited by the models.




Figure 4: Accuracy improvement over baseline for the best model with and without sentiment

Discussion and Concluding Remarks

In this study, we examined whether social sentiment holds predictive power regarding short-term stock returns. We scraped 28 months' worth of Tweets for the 10 most talked-about companies on Twitter by their cashtag. Two different sentiment analysis methods were applied to classify tweets as positive, negative, or neutral: VADER, a lexicon-based prebuilt model, and a custom-built sentiment analysis classifier built on a corpus of 3,000 manually labeled tweets. The sentiment scores were aggregated on a daily level. Subsequently, linear VARX and non-linear tree-based models were used to predict the following days' binarized return (positive or negative) from daily sentiment scores and financial indicators, putting forth three main findings:

1. The social media sentiment holds predictive power for 3 out of 10 companies in our sample.
2. Predictions based on the custom machine learning sentiment perform better than ones based on VADER sentiment.
3. Linear models deliver less accurate predictions than non-linear models.

Elaborating on finding 1, for Microsoft, Netflix, and Tesla sentiment holds predictive power, that is, there is a model using sentiment features that is not only better than the benchmark but also better than its counterpart without sentiment-based features. When adding sentiment features to the model, predictive accuracy increases for around 5%p to 9%p. This effect is smaller than the 18%p increase [32] find, but still cannot be attributed to random chance. For three other companies, Intel, AMD, and Nvidia, predictive models exist, but they do not use sentiment-based features. Finally, for the remaining four companies we were unable to find any predictive model with or without sentiment. This raises the question of whether some stocks possess inherent unpredictability, as all tested models failed for the same stocks. While previous research could not yet identify circumstances under which return prediction works particularly well, it has been shown that return forecasting only works for a minority of stocks:

Experiments conducted by [34] show that a significant relationship between sentiment and return exists only for 7% of the companies in their study. This also confirms statements by [39] who suggest that the value of sentiment for return forecasting needs to be examined on a case-by-case basis as no generalizable pattern exists.

Regarding the sentiment analysis technique (finding 2), for most of the companies (except for Tesla), the best model is produced by using the custom machine learning sentiment, not VADER. While the VARX models showed a contrary pattern where VADER seems to work slightly better than the machine learning sentiment, most of the best-performing models use the custom sentiment scores. This indicates that better sentiment assessment can lead to better forecasting performance. Therefore, especially considering the relatively bad performance of VADER on domain-specific texts, researchers should devote more resources towards the process of sentiment analysis and should carefully consider the use of off-the-shelf models trained on generic texts. Further evidence for this is provided in Appendix G, which displays the most predictive words for positive and negative sentiment in the custom model, most of which are domain-specific financial terms or social media slang. This specific vocabulary cannot be captured by generic sentiment models like VADER.

Finally, examining the performance of the VARX models shows that the magnitude of the performance increase is small with only several percentage points for all companies but Intel. While fitting the VARX model we find that for most companies, an order (number of lags included in the model) around two is optimal. This confirms findings from [40] that for day-ahead forecasting only a few days of lag are needed.

On the theoretical front, the large accuracy improvements over a naïve baseline for Microsoft, Netflix, and Tesla provide evidence against the Efficient Market Hypothesis under which no forecasting model should outperform random guessing. We show that consideration of the emotions of the crowd’s opinions yield better predictions of stocks dynamics, thus conceptually favoring BF over EMH. However, sentiment power is salient for a minority of stocks, which both confirms prior research [34] and might explain why some studies which only analyze a single or very few companies or indices conclude that sentiment cannot be used for predictive modeling. Our results warrant future research in the domains of sentiment analysis and stock return forecasting, implying larger samples need to be used to further examine circumstances that are conducive to predictable returns.

This study certainly comes with limitations. First, while assuming Twitter is a good proxy for public sentiment, there are other social networks and platforms on which users can discuss stock markets and share their opinions. As Twitter is a predominantly English-speaking platform, the sample of companies studied only includes large-cap companies from the United States. Results thus need not generalize to other markets, although similar studies have been conducted for Chinese markets [32][24] and cryptocurrency markets [26]. Moreover, social media sentiment can be operationalized in many ways. Here, we classify tweets into one of three classes and aggregate metrics on a daily level. Other approaches include measuring different types of emotions [22] or using network-based approaches. Finally, the results indicate that complex, non-linear forecasting models work better than simple ones. The application of Deep Learning to this data might thus yield further performance improvements.

References

1. Fama, E. (1970). Efficient Capital Markets: A Review of Theory and Empirical Work. *The Journal of Finance*, 25(2), 383-417. doi:10.2307/2325486
2. Shiller, R. J. (2003). From efficient markets theory to behavioral finance. *Journal of economic perspectives*, 17(1), 83-104.
3. Ritter, J. R. (2003). Behavioral finance. *Pacific-Basin finance journal*, 11(4), 429-437.
4. Brokernotes (2018). The modern trader. Retrieved from https://brokernotes.co/wpcontent/uploads/2017/08/BN-research-report_2018-FINAL.pdf.
5. Burton, K., Parmar, H. (January 27th, 2021). Reddit Crowd Bludgeons Melvin Capital in Warning to Industry. Bloomberg. Retrieved from <https://www.bloomberg.com/news/articles/2021-0127/bros-on-reddit-bludgeon-melvin-capital-in-warning-to-wall-street>
6. Walsh J. (2021). SEC Reportedly Warned Tesla To Stop Letting Elon Musk Tweet Without Permission. URL: <https://www.forbes.com/sites/joewalsh/2021/06/01/sec-reportedly-warned-tesla-to-stop-letting-elon-musk-tweet-without-permission/?sh=7ee5e5485afc>
7. Berger, J., & Milkman, K. L. "What makes online content viral?" *Journal of marketing research*, 2012, 49(2), 192-205.
8. Heimbach, I., & Hinz, O. "The impact of content sentiment and emotionality on content virality." *International Journal of Research in Marketing*, 2016, 33(3), 695-701.
9. Stieglitz, S., & Dang-Xuan, L. "Emotions and information diffusion in social media—sentiment of microblogs and sharing behavior." *Journal of management information systems*, 2013, 29(4), 217-248.
10. Xu, W. W., & Zhang, C. "Sentiment, richness, authority, and relevance model of information sharing during social Crises—the case of# MH370 tweets." *Computers in Human Behavior*, 2018, 89, 199-206.
11. Chen, L., Baird, A., & Straub, D. "A linguistic signaling model of social support exchange in online health communities." *Decision Support Systems*, 2020, 130, 113233.
12. Fama, E. F., Fisher, L., Jensen, M. C., & Roll, R. (1969). The adjustment of stock prices to new information. *International economic review*, 10(1), 1-21.
13. Lorie, J. H., & Niederhoffer, V. (1968). Predictive and statistical properties of insider trading. *The Journal of Law and Economics*, 11(1), 35-53.
14. Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *science*, 185(4157), 1124-1131.
15. De Bondt, W. F., Muradoglu, Y. G., Shefrin, H., & Staikouras, S. K. (2008). Behavioral finance: Quo Vadis?. *Journal of Applied Financ*), 18(2).
16. Glaser, M., Nöth, M., & Weber, M. (2004). Behavioral finance. *Blackwell handbook of judgment and decision making*, 527-546.
17. Tetlock, P. C., Saar-Tsechansky, M., & Macskassy, S. (2008). More than words: Quantifying language to measure firms' fundamentals. *The Journal of Finance*, 63(3), 1437-1467.
18. Khan, W., Ghazanfar, M. A., Azam, M. A., Karami, A., Alyoubi, K. H., & Alfakeeh, A. S. (2020). Stock market prediction using machine learning classifiers and social media, news. *Journal of Ambient Intelligence and Humanized Computing*, 1-24.
19. Yu, Y., Duan, W., & Cao, Q. (2013). The impact of social and conventional media on firm equity value: A sentiment analysis approach. *Decision support systems*, 55(4), 919-926.
20. Karabulut, Y. (2013, October). Can Facebook predict stock market activity?. In *AFA 2013 San Diego Meetings Paper*.
21. Siganos, A., Vagenas-Nanos, E., & Verwijmeren, P. (2014). Facebook's daily sentiment and international stock markets. *Journal of Economic Behavior & Organization*, 107, 730-743.

22. Bollen, J., Mao, H., & Zeng, X. (2011). Twitter mood predicts the stock market. *Journal of computational science*, 2(1), 1-8.
23. Mittal, A., & Goel, A. (2012). Stock prediction using Twitter sentiment analysis. Stanford University, CS229 (2011 <http://cs229.stanford.edu/proj2011/GoelMittal-StockMarketPredictionUsingTwitterSentimentAnalysis.pdf>), 15.
24. Li, Y., Bu, H., Li, J., & Wu, J. (2020). The role of text-extracted investor sentiment in Chinese stock price prediction with the enhancement of deep learning. *International Journal of Forecasting*, 36(4), 1541-1562.
25. Nti, I. K., Adekoya, A. F., & Weyori, B. A. (2020). Predicting stock market price movement using sentiment analysis: Evidence from Ghana. *Applied Computer Systems*, 25(1), 33-42.
26. Xie, P. (2021). The Interplay Between Investor Activity on Virtual Investment Community and the Trading Dynamics: Evidence From the Bitcoin Market. *Information Systems Frontiers*, 1-17.
27. Renault, T. (2019). Sentiment analysis and machine learning in finance: a comparison of methods and models on one million messages. *Digital Finance*, 1-13.
28. Audrino, F., Sigrist, F., & Ballinari, D. (2020). The impact of sentiment and attention measures on stock market volatility. *International Journal of Forecasting*, 36(2), 334-357.
29. Tabari, N., Biswas, P., Praneeth, B., Seyeditabari, A., Hadzikadic, M., & Zadrozny, W. (2018, July). Causality analysis of Twitter sentiments and stock market returns. In *Proceedings of the first workshop on economics and natural language processing* (pp. 11-19).
30. Rao, T., & Srivastava, S. (2014). Twitter sentiment analysis: How to hedge your bets in the stock markets. In *State of the art applications of social network analysis* (pp. 227-247). Springer, Cham.
31. Teti, E., Dallochio, M., & Aniasi, A. (2019). The relationship between Twitter and stock prices. Evidence from the US technology industry. *Technological Forecasting and Social Change*, 149, 119747.
32. Ren, R., Wu, D. D., & Liu, T. (2018). Forecasting stock market movement direction using sentiment analysis and support vector machine. *IEEE Systems Journal*, 13(1), 760-770.
33. Sun, T., Wang, J., Zhang, P., Cao, Y., Liu, B., & Wang, D. (2017, August). Predicting stock price returns using microblog sentiment for chinese stock market. In *2017 3rd International Conference on Big Data Computing and Communications (BIGCOM)* (pp. 87-96). IEEE.
34. Coqueret, G. (2020). Stock-specific sentiment and return predictability. *Quantitative Finance*, 20(9), 1531-1551.
35. 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(3), 1259-1294.
36. Oliveira, N., Cortez, P., & Areal, N. (2017). The impact of microblogging data for stock market prediction: Using Twitter to predict returns, volatility, trading volume, and survey sentiment indices. *Expert Systems with Applications*, 73, 125-144.
37. Twitter Inc. (2019). Selected Company Financials and Metrics. Retrieved from https://s22.q4cdn.com/826641620/files/doc_financials/2019/q1/Q1-2019-Selected-Company-Metrics-and-Financials.pdf (access: June 14th, 2021).
38. Twintproject (2021). TWINT – Twitter Intelligence Tool. Retrieved from <https://github.com/twintproject/twint.git> (access: May 6th, 2021).
39. Mudinas, A., Zhang, D., & Levene, M. (2019). Market trend prediction using sentiment analysis: lessons learned and paths forward. *arXiv preprint arXiv:1903.05440*.
40. Shynkevich, Y., McGinnity, T. M., Coleman, S. A., Belatreche, A., & Li, Y. (2017). Forecasting price movements using technical indicators: Investigating the impact of varying input window length. *Neurocomputing*, 264, 71-88.