

RESEARCH ARTICLE

Social media engagement and cryptocurrency performance

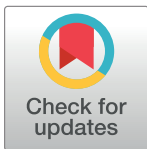
Khizar Qureshi ^{*}, Tauhid Zaman

Yale School of Management, Yale University, New Haven, CT, United States of America

^{*} khizar.qureshi@yale.edu

Abstract

Cryptocurrencies are highly speculative assets with large price volatility. If one could forecast their behavior, this would make them more attractive to investors. In this work we study the problem of predicting the future performance of cryptocurrencies using social media data. We propose a new model to measure the engagement of users with topics discussed on social media based on interactions with social media posts. This model overcomes the limitations of previous volume and sentiment based approaches. We use this model to estimate engagement coefficients for 48 cryptocurrencies created between 2019 and 2021 using data from Twitter from the first month of the cryptocurrencies' existence. We find that the future returns of the cryptocurrencies are dependent on the engagement coefficients. Cryptocurrencies whose engagement coefficients have extreme values have lower returns. Low engagement coefficients signal a lack of interest, while high engagement coefficients signal artificial activity which is likely from automated accounts known as bots. We measure the amount of bot posts for the cryptocurrencies and find that generally, cryptocurrencies with more bot posts have lower future returns. While future returns are dependent on both the bot activity and engagement coefficient, the dependence is strongest for the engagement coefficient, especially for short-term returns. We show that simple investment strategies which select cryptocurrencies with engagement coefficients exceeding a fixed threshold perform well for holding times of a few months.



OPEN ACCESS

Citation: Qureshi K, Zaman T (2023) Social media engagement and cryptocurrency performance. PLoS ONE 18(5): e0284501. <https://doi.org/10.1371/journal.pone.0284501>

Editor: Aurelio F. Bariviera, URV: Universitat Rovira i Virgili, SPAIN

Received: August 11, 2022

Accepted: April 2, 2023

Published: May 11, 2023

Copyright: © 2023 Qureshi, Zaman. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: "All of the supporting source code and data to replicate the results are included in a public GitHub repository (https://github.com/kai-trading-bot/crypto_engagement)."

Funding: The authors received no specific funding for this work.

Competing interests: The authors have declared that no competing interests exist.

Introduction

The cryptocurrency market has grown rapidly since the creation of bitcoin in 2008. Initially the market was comprised of a single cryptocurrency, but by 2022 there were over 10,000 active cryptocurrencies with a combined market capitalization near 2 trillion USD [1]. Cryptocurrencies have become a popular speculative investment, leading to a proliferation of exchanges where cryptocurrencies are actively traded 24 hours a day, seven days a week [2]. Unlike more traditional asset classes, cryptocurrency prices are not primarily driven by concrete measures such as earnings. Cryptocurrency markets have been found to react to economic surprise shocks like traditional asset classes [3], and studies have also found that the cryptocurrency market is dominated by irrational investors who base their investment

decisions on market sentiment [4]. This leads one to ask to what extent can future cryptocurrency returns be predicted by sentiment data and how does one go about making such predictions.

Cryptocurrency markets are highly volatile and can have extreme reactions to external events. It has been found that local and global shocks affect local bitcoin activities and trading volatilities, confirming that economic events affect bitcoin markets. This is especially true of cryptocurrencies with small market capitalizations, commonly referred to as *alt-coins*. One famous example of an alt-coin reacting to external events is dogecoin. This is an alt-coin initially created as a way to mock the speculative nature of cryptocurrencies. Nevertheless, it was able to reach a market capitalization of 50 billion USD in 2021 [5]. Elon Musk, the founder of Tesla Motors, took an interest in dogecoin and on Twitter he would often post humorous messages, known as tweets, about it. The price of dogecoin would often fluctuate wildly in response to Musk's tweets [6].

While dogecoin is an extreme case, it does illustrate the potential for social media to predict the future performance of cryptocurrencies, especially alt-coins. Social media can capture the sentiment and general level of engagement of a population towards a variety of topics. Several studies have attempted to measure user engagement in online social networks [7–10]. However, these approaches are all limited to either surveys or “black-box” empirical methods, both of which carry their own limitations. These studies all assume that more clicks or general interactions with social media posts are indicative of higher engagement.

For cryptocurrencies, in principle a positive sentiment or high level of engagement can lead to positive future returns. In [11] it was found that cryptocurrency prices react positively to Twitter sentiments of the issuers of the cryptocurrencies. This was demonstrated in [12] where it was found that short-term returns (24 hours) of alt-coins could be predicted using activity and sentiment from Twitter. A trading strategy for bitcoin utilizing social media data, in addition to many other signals is presented in [13]. The social signals here are rather simple, consisting of number of posts about the cryptocurrency and the mean emotional valence of the posts. In [14] the authors construct a measure for economic policy uncertainty using Twitter data, and then show the impact on several cryptocurrencies with large market capitalization.

In the context of social media manipulation, it was found that alt-coins were often exploited by coordinated “pump and dump” schemes orchestrated by fraudulent groups [15]. These are schemes where artificial social media activity is used to create excitement for an alt-coin and pump up the price, after which the perpetrators of the scheme sell the alt-coin before the price dumps down to extremely low levels. Whether genuine or artificial, it does appear that social media activity can tell us something about the future performance of a cryptocurrency. For example, [11] find that social sentiment has a positive correlation with cryptocurrency and that volume is reactive to the absolute value in social sentiment. Additionally, studies have found causal relationships between Twitter-based measures of economic and market uncertainty and prices for many major cryptocurrencies [14, 16].

The power of social media to predict financial performance has been demonstrated outside of cryptocurrencies. It has been shown that the volume of social media posts about a film are predictive of its ticket sales [17]. In traditional finance, there are many works that use social media data to predict the prices of stocks [18–26]. These works often utilize the sentiment of social media posts in addition to their volume. While social media can provide a measure of sentiment, there are other approaches that measure investor sentiment for equity indices [27]. These have traditionally been calculated using consumer and investment management surveys. One limitation of these measures is that they are calculated periodically, such as every quarter, rather than in real-time. A comprehensive overview of studies on investor behavior can be found in [4].

Many approaches attempting to predict financial performance using social media rely upon the volume as a predictive feature. This is reasonable if one has access to all relevant social media posts. However, many times one faces data limitations and only has access to a sample of the posts. In this case the volume feature will be censored and one cannot gauge its true value. Even if one has access to all posts, the data demands of the volume feature can be high. The typical daily number of tweets about a cryptocurrency can vary from tens for less-known cryptocurrencies to hundreds of thousands for some of the more common ones. To construct the volume feature, an excessive number of posts must be collected for popular cryptocurrencies. If one reached a data collection limit, one might mistakenly conclude that the volume was equivalent for extremely popular and less popular cryptocurrencies.

Sentiment is another dimension of social media data used to predict financial performance. Conventional approaches analyze the text to measure the sentiment of a social media post. However, the language used on social media can evolve rapidly and may not be understood by sentiment analysis tools. This can include hashtags created by users in the online cryptocurrency community, such as #buythedip or #hodl, both of which express positive sentiment, but may not be comprehensible to sentiment analysis tools which are not trained on cryptocurrency social media data. Other times, social media posts contain no text at all and instead are just images. These “memes” as they are commonly known, contain strong sentiment signals which cannot be extracted by text analysis.

Instead of volume or text based sentiment, a more useful social media signal can be obtained from the level of engagement from the followers of social media users. Many social media platforms have network structures, where users have followers who are able to see their posts. These followers can then show their engagement with the topic of the post by interacting with it in different ways. For instance, in Twitter, one interaction is called liking, where a user clicks on a post to signal that they like it. Interaction counts are a good way to gauge the engagement of users with a topic on social media. First, they are language agnostic. If a post is simply an image, or a cryptic hashtag, the interaction count can be measured. Second, they are not subject to data limitations. If a post receives an incredibly large number of interactions, this value can be easily measured, no matter how large. This is in contrast to social media volume, where one must collect all posts about a topic.

In this work we propose a generative model for social media engagement for different topics based on the interaction counts of relevant social media posts. A key element of this model is the *engagement coefficient* for a topic, which measures how engaged social media users are with posts about the topic. This model is language agnostic and not vulnerable to data collection limitations. We derive closed form expressions for the engagement parameter maximum likelihood estimates. We then apply this model to cryptocurrencies. We collect price histories and social media posts for all alt-coins created between 2019 and 2022 with sufficiently high fund-raising goals. In total our dataset contains over 1.36 million tweets for 48 different alt-coins. We estimate the engagement coefficient for these alt-coins and analyze their relationship with future returns. We find that returns are low if the engagement coefficient is too low or too high. Too low indicates lack of engagement, while too high suggests artificial engagement from artificial accounts, known as bots. We analyze the prevalence of bot posts for each cryptocurrency and find that a higher prevalence is associated with lower future returns. We find that investment strategies based on the engagement coefficients can achieve significant short-term returns (one month), but are less effective for long-term returns (one year). We provide the dataset used in this study in the Supporting Information ([S1 Data](#)).

In the extant literature, the most relevant works are [12, 13], both of which attempt to measure the relationship between cryptocurrency and social media data. However, our work has many important differences. First, both of these works use simple signals constructed from the

social media data, such as volume and sentiment. In contrast, we explicitly calculate a novel measure of engagement based on social media data in a manner that is not vulnerable to data collection limits nor requires any type of text-based sentiment measurement algorithm. Second, in both of these works the time-horizon of prediction of cryptocurrency prices using social media data is very short (hours or days). Our approach develops a social media signal that can be utilized for longer-term investment decisions (weeks or months). Third, these works consider cryptocurrencies that have been in existence for a significant amount of time. Our work focuses on cryptocurrencies during the first month of their existence, and thus studies the question of deciding whether or not to invest in a newly created cryptocurrency.

Materials and methods

Social media engagement model

We consider a topic or item c that has some quantifiable performance metric. For instance, c could be a movie and the metric is ticket sales, or c could be a politician and the metric is votes received in an election. In this work we will assume c is a cryptocurrency and the performance metric is future returns. We assume that if there is a great deal of excitement and enthusiasm for the topic, then in the future its performance metric will increase. For instance, if people are excited about a new cryptocurrency, this will lead them purchase it, which will increase its price and future returns. We capture this excitement for the topic with a latent engagement coefficient α_c . If the current excitement for the topic affects the performance metric in the future, then one can in principle use the engagement coefficient to predict these future values.

To predict future performance using the engagement coefficient, we need data where people manifest their excitement for a topic in some measurable manner before the future performance metrics are realized. We can then use this data to estimate α_c , and in turn use α_c to predict future performance. Social media platforms allow individuals to show such excitement. For instance, consider Twitter, a large social media platform. A Twitter user can post a tweet discussing a topic. Once this tweet is posted, it becomes visible to the followers of this user. The followers can show their excitement for the topic by interacting with the tweet. There are three different ways Twitter users can interact with a tweet. First, they can *like* the tweet by clicking the like button on the tweet. Second, they can *retweet* the tweet by clicking the retweet button on the retweet. Third, they can post a *reply* to the original tweet. Liking will only change the like count visible on the tweet. Retweeting will post the original tweet on the retweeting user's Twitter timeline, making it visible to their own followers. Replies are tweets, and so will also be visible to the replying user's followers.

One important factor to consider with interaction counts is the follower count of the posting users. Generally, users with more followers will receive more interactions, as the bulk of these come from followers. Therefore, interaction counts are more useful if they are normalized by the follower count of the posting user. Then one can view engagement with a topic as the fraction of followers who interact with posts about it. This way we can define a uniform measure of engagement which is not sensitive to follower count. Normalizing by follower count has been used in previous works trying to forecast the reach of social media posts on Twitter [28].

We assume that the number of interactions a tweet receives will be higher if the tweet is about a topic with a larger engagement coefficient. Naturally, if people are more excited about a topic, this should increase the probability that they will interact with a tweet about the topic. Therefore, we can use social media data on interactions with tweets to estimate engagement coefficients for different topics.

Model specification

We now present a generative model for social media interactions about a topic based on its latent engagement coefficient. For simplicity, we assume the social media platform is Twitter. However, the model can easily be generalized to other social media platforms with different interaction types.

We consider a set \mathcal{C} of topics. For each topic $c \in \mathcal{C}$, we associate an engagement coefficient $\alpha_c \geq 0$. We observe a set \mathcal{U} of social media users. Each user $u \in \mathcal{U}$ has a follower count $f_u \geq 0$. These followers are able to see tweets posted by u and interact with them. We assume there is a discrete set \mathcal{I} of interactions that can be performed on a tweet. For each interaction $i \in \mathcal{I}$ we associate an interaction coefficient $\beta_i \geq 0$ which quantifies the effort needed to perform the interaction. For instance, in Twitter, liking has a relatively low effort, as it only requires clicking a button.

Each user u has m_{cu} tweets about topic c . Tweet k by user u for topic c has $n_{cui k}$ type i interactions. We assume that $n_{cui k}$ is a Poisson random variable with mean μ_{cui} given by

$$\mu_{cui} = \beta_i \alpha_c f_u$$

Our specification assumes the mean number of interactions scales linearly with the follower count of the user. Generally, the interactions on a user’s tweet come from their followers, and users with more followers have higher mean interaction counts. From our specification we also see that a higher engagement coefficient and a higher interaction coefficient will lead to a higher mean interaction count.

Our dataset consists of interaction counts for a set of tweets from multiple users about multiple topics. We denote this dataset as $\mathbf{n} = \{n_{cui k}\}_{k=1, c \in \mathcal{C}, u \in \mathcal{U}, i \in \mathcal{I}}$. We define the set of engagement coefficients as $\alpha = \{\alpha_c\}_{c \in \mathcal{C}}$ and the set of interaction coefficients as $\beta = \{\beta_i\}_{i \in \mathcal{I}}$. With this notation, the log-likelihood of the dataset is

$$L(\alpha, \beta | \mathbf{n}) = \sum_{c \in \mathcal{C}} \sum_{u \in \mathcal{U}} \sum_{i \in \mathcal{I}} \sum_{k=1}^{m_{cu}} (-\mu_{cui} + n_{cui k} \log(\mu_{cui}) - \log(n_{cui k}!)).$$

To simplify notation, we make the following definitions for the total interactions and the total interactions for a topic, a (user, topic) pair, a (user,topic,interaction) triplet, and a single interaction:

$$n = \sum_{c \in \mathcal{C}} \sum_{i \in \mathcal{I}} \sum_{u \in \mathcal{U}} \sum_{k=1}^{m_{cu}} n_{cui k}$$

$$n_c = \sum_{i \in \mathcal{I}} \sum_{u \in \mathcal{U}} \sum_{k=1}^{m_{cu}} n_{cui k}$$

$$n_{cu} = \sum_{i \in \mathcal{I}} \sum_{k=1}^{m_{cu}} n_{cui k}$$

$$n_{cui} = \sum_{k=1}^{m_{cu}} n_{cui k}$$

$$l_i = \sum_{c \in \mathcal{C}} \sum_{u \in \mathcal{U}} \sum_{k=1}^{m_{cu}} n_{cui k}$$

We also define the maximum number of follower interactions for a single interaction type for the tweets posted about a topic c as

$$v_c = \sum_{u \in \mathcal{U}} m_{cu} f_u.$$

For a single interaction type, if all followers of the users in \mathcal{U} performed this interaction with all of their tweets about topic c , the total number of interactions would be v_c .

With this notation we can rewrite the log-likelihood as

$$L(\alpha, \beta | \mathbf{n}) = - \sum_{c \in \mathcal{C}} \alpha_c v_c \left(\sum_{i \in \mathcal{I}} \beta_i \right) + \sum_{c \in \mathcal{C}} n_c \log(\alpha_c) + \sum_{i \in \mathcal{I}} l_i \log(\beta_i) + C.$$

Above, we put all terms not dependent on the model parameters into the constant C .

To fit the model we use maximum likelihood (ML) estimation. The gradients of the log-likelihood for each parameter are set to zero and then this system of equations is solved to obtain the model parameters. However, in the model's current form, there are multiple solutions. For instance, we can multiply the α_c 's by a positive constant and divide the β_i 's by the same constant, and have an equivalent model. To obtain a unique solution, we fix the value of one of the interaction coefficients. Without loss of generality, we set $\beta_1 = 1$ for interaction type one. Then, we obtain the following closed form solutions for the parameter ML estimates, with a detailed derivation provided in the Supplement Information ([S1 Appendix](#)):

$$\hat{\alpha}_c = \frac{n_c n}{v_c l_1}, \quad c \in \mathcal{C} \tag{1}$$

$$\hat{\beta}_i = \frac{l_i}{l_1}, \quad i \in \mathcal{I}. \tag{2}$$

If we consider the case of a single interaction type, the ML solution becomes very easy to interpret. To obtain a unique solution, we do not include any interaction coefficients and set $\mu_{cu} = \alpha_c f_u$. The ML solution is

$$\hat{\alpha}_c = \frac{n_c}{v_c}, \quad c \in \mathcal{C}.$$

We see from this solution that the engagement coefficient is simply the ratio of the number of interactions divided by the maximum possible follower interactions for the topic. Essentially, it captures the fraction of followers that interact with posts about the topic. If all interactions only came from followers, the engagement coefficient would have a maximum value of one. In reality, non-followers can also interact with the posts, so the value can exceed one. However, we will see that typically the value is much less than one.

Performance data

In this study we consider cryptocurrencies created between 2019 and 2021. There are different ways that cryptocurrencies are created and initially offered for sale. An initial coin offering (ICO) involves a company's sale of a coin, app, or service in order to raise funds for the continuation of the underlying project. In most cases, the company sets a fixed fundraising goal and a fixed supply of the coin. More recent variations include a dynamic price and supply based on the public's interest. An initial exchange offering (IEO), on the other hand, involves launching the coin directly on a centralized exchange. The projects are carefully examined and vetted before the launch, reducing the risk of scam and fraud often found in ICOs. Finally, an initial

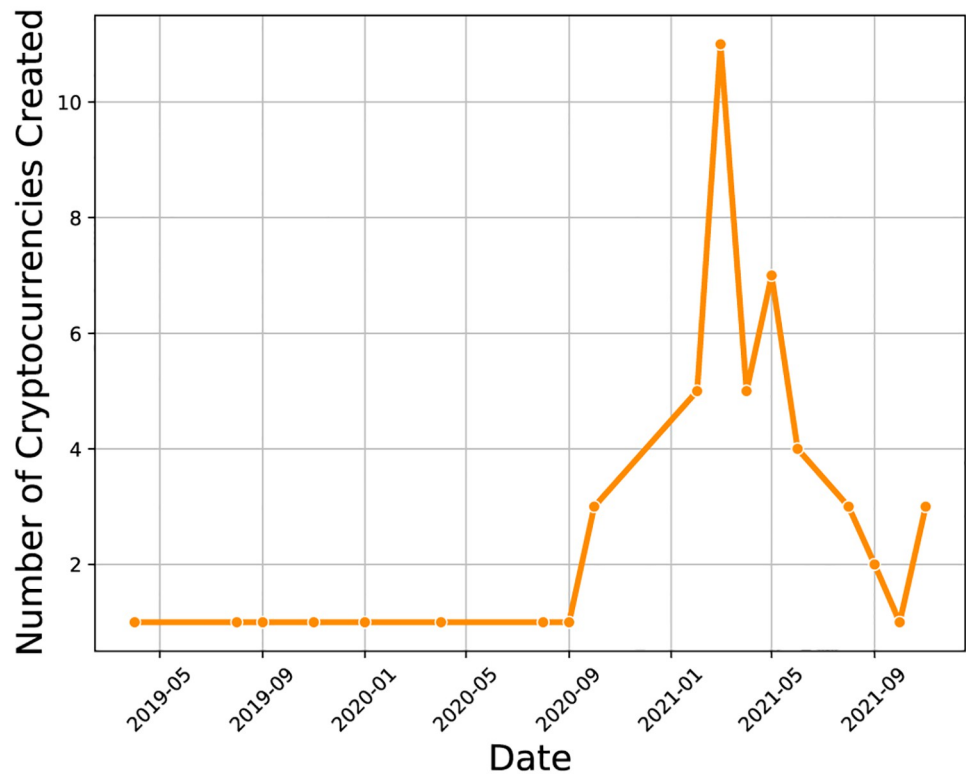


Fig 1. The number of cryptocurrencies in our dataset created each month between 2019 and 2021.

<https://doi.org/10.1371/journal.pone.0284501.g001>

decentralized exchange offering (IDO) involves the listing on a decentralized exchange instead of a centralized one. IDOs also differ from IEOs since they tend to have a lower review time as well as lower listing costs.

For this study we considered cryptocurrencies created between 2019 and 2021. The creation date of a cryptocurrency will refer to its initial offering date. We only considered cryptocurrencies with an ICO, IEO, or IDO whose fundraising goal was 1 million USD or higher. We used this fundraising threshold in order to focus on coins which could potentially have large market capitalizations. We fetched the initial offering data from the website CoinMarketCap [29]. Fig 1 shows the number of coins created each month. We see that the majority of the coins were created in 2021. In total we have 48 coins in the 2019–2021 time period. Our study is based on data from 2019–2021 also because this is when cryptocurrencies had substantial fundraising goals (greater than 1 million USD). We provide a complete list of the cryptocurrencies and their initial offering dates in the Supporting Information (S1 Table). For each of the cryptocurrencies meeting our fundraising threshold, we also collected the daily price data between the initial offering date to one year after from CoinGecko [30]. The performance metric for the cryptocurrencies that we consider is their price returns.

Social media data

We collected social media data about the cryptocurrencies from Twitter. For each cryptocurrency, we used a historical search via the Twitter API [31] to collect all tweets containing the *powertag* of the cryptocurrency that were posted between its creation date to one month after. A *powertag* is a dollar sign followed by the ticker symbol for the cryptocurrency. Powertags

are a common way people mention cryptocurrencies on Twitter. Examples of power tags include \$BTC for bitcoin, \$ETH for ethereum, and \$CAKE for cake. We chose one month as the time period for which we collect social media data because this is a reasonable amount of time to gauge the engagement of users with the cryptocurrency. In practice, one would estimate the engagement parameter with this one month of data, and then make an investment decision about the cryptocurrency.

In addition to tweets, we also collected the Twitter profiles for each user in our dataset. From the profiles we were able to obtain the follower count of the users, which is necessary to fit our model for the engagement coefficient. In total, our dataset contained 1.36 million tweets from 129,071 unique users. Detailed statistics of the social media dataset are provided in the Supporting Information (S1 Table).

Data statement

To the best of our knowledge, our methods for data collection were compliant with the relevant vendors. For example, our social media data was collected exclusively using the Twitter academic API, for which we thoroughly detailed our use case. Our market data was collected exclusively from CoinMarketCap.com and CoinGecko.com, both of which offers price data per coin for download publicly. The data used in this study can be found at <https://github.com/kai-trading-bot/engagement>.

Results

For each cryptocurrency, we have social media posts from the first month of its existence. The cryptocurrencies were created at different points in time, which allows for different ways to estimate the engagement model. First, we do a completely in-sample estimation using all of the social media data. This means that engagement coefficients for some cryptocurrencies will be estimated using future data. Second, we estimate the model separately for each cryptocurrency using only data available when it has existed for one month. Estimating the model in this manner, which we refer to as prior data estimation, avoids any look-ahead bias. In practice, an investor would estimate the model in this manner. After estimating the model, we analyze the relationship between the engagement coefficients of the cryptocurrencies and their future returns.

Model estimation

We apply Eq (1) to all the Twitter data collected for the cryptocurrencies to obtain the in-sample model parameter ML estimates. The resulting engagement coefficients are shown in Fig 2 and the interaction coefficients are listed in Table 1. From the figure we see that the engagement coefficients span several orders of magnitude, ranging from 5×10^{-6} up to 1.7×10^{-2} , and 43 of the 48 cryptocurrencies have engagement parameters in the narrower 10^{-3} to 10^{-5} range. There are two cryptocurrencies with engagement parameters greater than 10^{-3} : *krypto*, and *safebtc*. The lower range has *mpt*, *palg*, and *cspr* with engagement coefficients less than 10^{-5} .

We plot the engagement coefficients of the cryptocurrencies versus their creation date in Fig 3 to show an interesting temporal relationship. Cryptocurrencies created prior to 2021 have engagement coefficients almost entirely in the 10^{-4} to 10^{-3} range. In contrast, after 2021, the engagement coefficients show a larger variation, with many cryptocurrencies in the lower range of 10^{-5} to 10^{-4} . The majority of the cryptocurrencies were created in 2021, and from this we see that many of them had low engagement. Also, in 2021 we see two cryptocurrencies

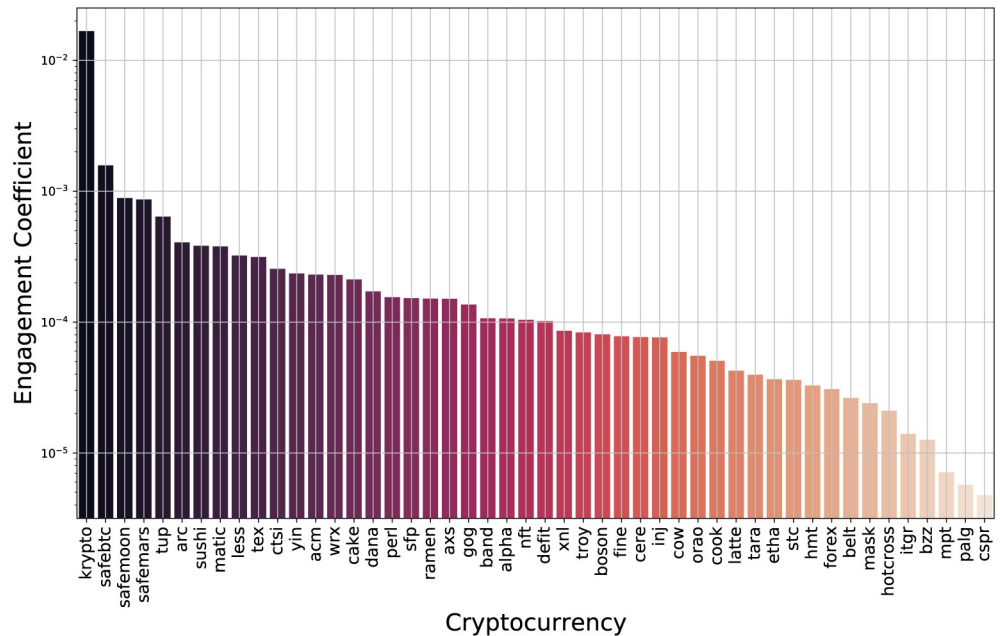


Fig 2. Plot of the in-sample estimated engagement coefficients.

<https://doi.org/10.1371/journal.pone.0284501.g002>

(krypto and safebtc) with engagement coefficients greater than 10^{-3} , and three below 10^{-5} (mpt, palg, and cspr).

The estimated interaction coefficients are shown in Table 1. Our model specification fixes the liking coefficient to one. The retweet coefficient is 0.31, indicating that retweets are one third as common as likes. Replying has an interaction coefficient of 0.19, making it the least common interaction. These values align with what is expected, as retweeting and replying require more effort than liking, an liking requires the most effort because one must write a reply tweet.

If one were to use the engagement coefficients to make investment decisions for cryptocurrencies, then data from the future could not be used to estimate the coefficients. To simulate this scenario, we estimate the engagement coefficients for each cryptocurrency using only data available one month after the creation date. This includes data for the cryptocurrency in question along with any other cryptocurrencies created before it. In this way, the estimated engagement coefficients can be used for investment decisions. This approach has a separate estimation performed for each cryptocurrency. We find that the engagement coefficients for each method are nearly equal. The mean absolute percent error between the engagement coefficients estimated using both methods is 1.3%, with a maximum absolute percent error of

Table 1. Interaction coefficients estimated using all data. Note that the interaction coefficient for liking is set equal to one.

Interaction	Interaction coefficient
Liking	1.00
Retweeting	0.31
Replying	0.19

<https://doi.org/10.1371/journal.pone.0284501.t001>

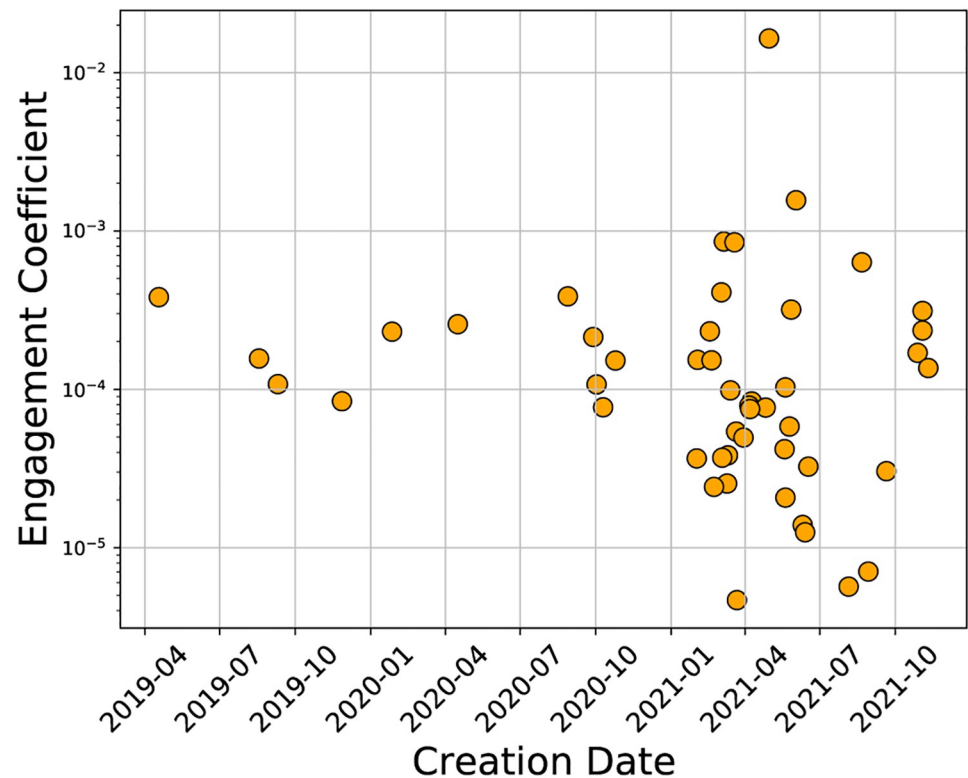


Fig 3. Plot of the cryptocurrency engagement coefficients versus creation date.

<https://doi.org/10.1371/journal.pone.0284501.g003>

3.4%. Therefore, we see that the model estimates for the engagement coefficients are robust to the data used.

Engagement coefficient and returns

Now that we have estimated the engagement coefficients using prior social media data, we can investigate their relationship to the future returns of the cryptocurrencies. We consider returns measured with respect to one month after the cryptocurrencies were created because the first month is an observation period during which we collect social media data to estimate the engagement coefficient.

We begin by looking at the cryptocurrencies with the top returns in for different holding times in [Table 2](#). We see that the returns become quite large, and the maximum return increases with time. After 12 months, there are cryptocurrencies with returns above 2,000%. The highest return is 24,985% for *axs*. This cryptocurrency is the governance token for the online video game *Axie Infinity* which experienced an incredible increase in popularity during that time period [[32](#)]. We note that for these top performing cryptocurrencies, the engagement coefficient has an order of magnitude of 10^{-4} , with the exception of *inj* which has a value of 0.77×10^{-5} .

To further visualize the relationship between the engagement coefficient and returns, we plot the evolution of the cryptocurrency returns over time as a function of the engagement coefficient in [Fig 4](#). We see that the cryptocurrencies with the highest returns have engagement coefficients between 10^{-4} and 10^{-3} . The returns are very low for cryptocurrencies with engagement coefficients outside of this range.

Table 2. Engagement coefficients and returns for cryptocurrencies with top returns in different time periods.

Cryptocurrency	Creation Date	Engagement Coeff.	1 Month Return [%]
safemoon	3/6/2021	8.58×10^{-4}	1,182
safemars	3/19/2021	8.48×10^{-4}	545
alpha	10/2/2020	1.07×10^{-4}	361
Cryptocurrency	Creation Date	Engagement Coeff.	3 Month Return [%]
alpha	10/2/2020	1.07×10^{-4}	4,312
inj	10/10/2020	7.7×10^{-5}	952
axs	10/25/2020	1.52×10^{-4}	328
Cryptocurrency	Creation Date	Engagement Coeff.	6 Month Return [%]
cake	9/28/2020	2.14×10^{-4}	8,829
alpha	10/2/2020	1.07×10^{-4}	3,843
sushi	8/28/2020	3.87×10^{-4}	1,078
Cryptocurrency	Creation Date	Engagement Coeff.	9 Month Return [%]
axs	10/25/2020	1.52×10^{-4}	14,329
cake	9/28/2020	2.14×10^{-4}	3,587
alpha	10/2/2020	1.07×10^{-4}	1,902
Cryptocurrency	Creation Date	Engagement Coeff.	12 Month Return [%]
axs	10/25/2020	1.52×10^{-4}	24,985
cake	9/28/2020	2.14×10^{-4}	4,409
alpha	10/2/2020	1.07×10^{-4}	2,416

<https://doi.org/10.1371/journal.pone.0284501.t002>

Very low engagement coefficients are due to the followers not interacting with posts about the cryptocurrency. This signals a lack of interest, which we expect leads to low returns. The surprising aspect is that extremely high engagement coefficients also have low returns. We would expect that if the followers are likely to interact with posts about a cryptocurrency, then there is strong interest in it, leading to a price increase. However, we find this is not the case. We suspect that the reason we do not see this monotonic type of behavior is due to artificial Twitter accounts, known as bots. Studies have suggested that between 9% and 15% of active Twitter accounts are bots [33]. One can create a large number of bots and have them interact with posts about a cryptocurrency to create the impression that there is excitement for it, which is typical of pump and dump schemes [15]. This can lead to engagement coefficients being unusually high. From Fig 2 we see that the cryptocurrency with the highest engagement coefficient is krypto, which is used for in-game purchases in the online game Kryptobellion. The returns for krypto were as high as 49% after two months, but then dropped to negative values thereafter. The engagement coefficient for krypto is 1.67×10^{-2} , which is nearly an order of magnitude larger than the next highest value. This is likely due to bots programmed to like and retweet any tweets about krypto. However, with the data we have available, we cannot make this claim with certainty. Nonetheless, the engagement coefficient can still provide some indication that there is artificial manipulation of social media activity. Given our analysis here, we would suggest 10^{-3} as a threshold for when there is possible manipulation.

Bot tweets and returns

Our analysis of the engagement coefficients suggested that there may be bots artificially increasing engagement for different cryptocurrencies by interacting with tweets. The presence of bots tweeting about a cryptocurrency may be indicative of future performance. For instance,

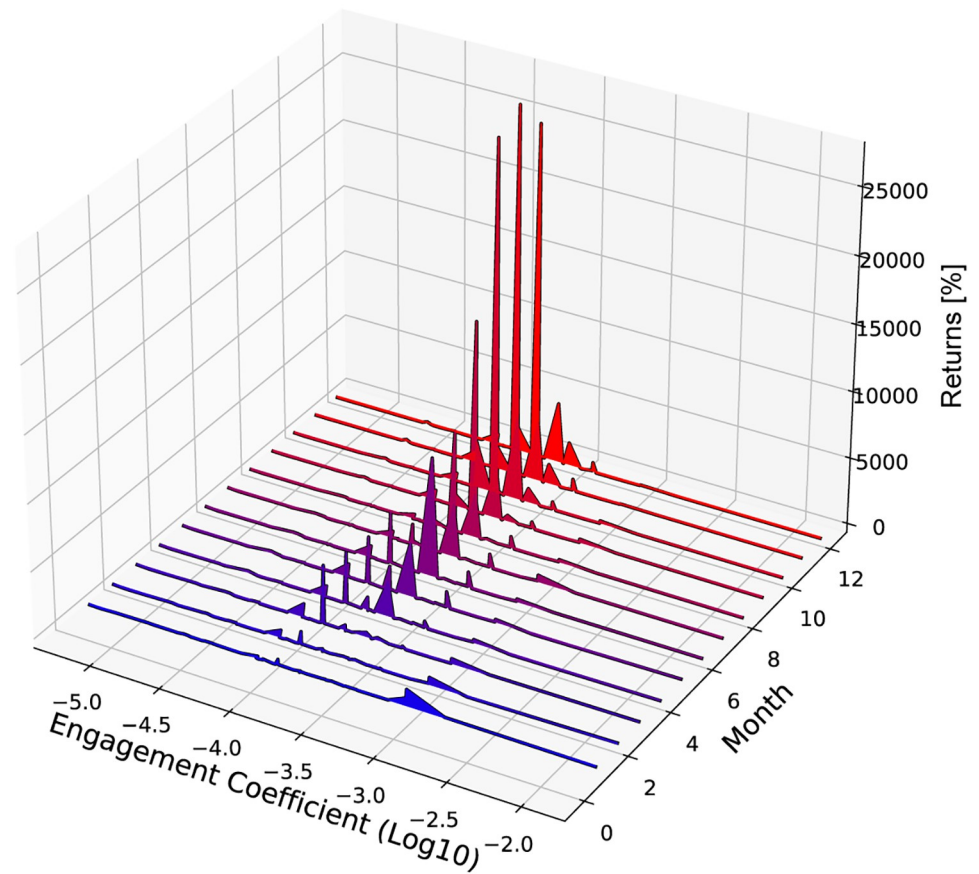


Fig 4. Plot of the evolution of the returns of different cryptocurrencies as a function of the logarithm of their engagement coefficient (estimated using prior data). The returns are measured relative to one month after the cryptocurrency was created.

<https://doi.org/10.1371/journal.pone.0284501.g004>

a cryptocurrency whose social media discussion contains a large fraction of bots may be a pump and dump scheme.

Unfortunately, we are not able to determine which tweet interactions come from bots. However, in addition to interacting with tweets, bots can also post tweets about cryptocurrencies, and unlike interactions, it is possible to identify which tweets come from bots. To do this, we used the Botometer, a machine learning algorithm that calculates the probability that a Twitter account is a bot based on its posted tweets, profile information, and local follower network [34]. The Botometer has been used in other studies to identify bot accounts in Twitter [35]. Although it is imperfect, as discussed in [36], its use is widely accepted in the academic community.

We used the Botometer API [37] to obtain the bot probability for each user in our dataset. The Botometer bot probabilities can be used to quantify the amount of bot discussion for a cryptocurrency. One measure of bot prevalence is the mean bot probability for all users discussing the cryptocurrency. We calculate this quantity for each cryptocurrency and show a histogram of the resulting values in Fig 5. The lowest mean bot probabilities are for sushi (0.25) and cake (0.31). Both of these are decentralized finance (DeFi) cryptocurrencies which achieved some success. The 12 month returns were 630% for sushi and 4,400% for cake. The highest mean bot probability belonged to latte (0.49), another DeFi cryptocurrency, but which

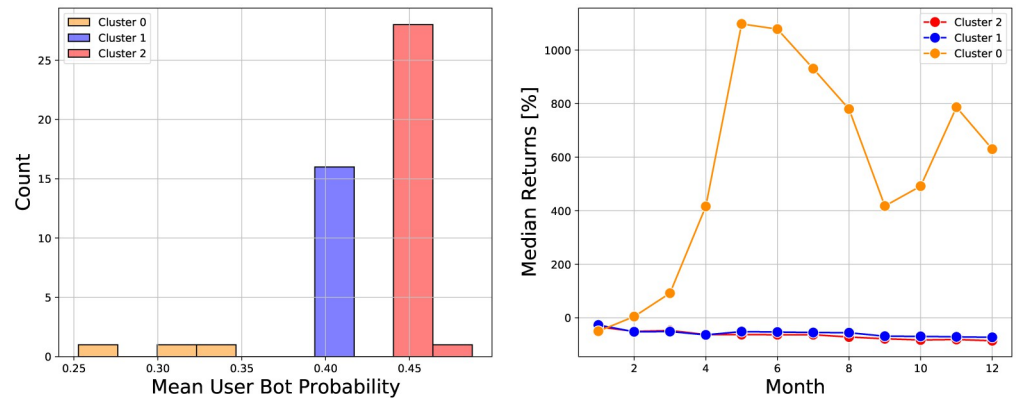


Fig 5. (left) Histogram of the cryptocurrency mean bot probabilities. (right) Plot of the median returns in each bot probability cluster versus time (relative to one month after the cryptocurrency's creation).

<https://doi.org/10.1371/journal.pone.0284501.g005>

did not achieve any success. The latte price decreased by 95% in three months and did not recover by the 12th month.

From visual inspection we find three clusters of mean bot probabilities. We apply k-means with three clusters to the mean bot probabilities and find clusters centered at 0.30, 0.40, and 0.44. The cluster at 0.30 has three cryptocurrencies, the cluster at 0.40 has 16 cryptocurrencies, and the cluster at 0.44 has 29 cryptocurrencies. To understand the relationship between the mean bot probability and performance, we plot the median returns of cryptocurrencies in the three different bot probability clusters versus time in Fig 5. We see that the later returns are highest for the cluster with the lowest mean bot probability. The two clusters with the highest bot probabilities show similar returns which are quite low.

Comparing the dependence of returns on social media features

We now have multiple social media features for cryptocurrencies. There is the engagement coefficient, which captures user enthusiasm. There is the mean user bot probability, which captures the prevalence of bots in the social media activity of a cryptocurrency. Finally, as a baseline feature, there is the total number, or volume, of tweets posted about a cryptocurrency in its first month. The tweet volume is related to the engagement coefficient because it attempts to measure enthusiasm. However, unlike the engagement coefficient, tweet volume is unbounded and does not take into account the follower counts of users.

Previous studies have used tweet volume to predict short-term future returns of cryptocurrencies [12]. We wish to see if our new social media features have any advantage over the volume feature. To compare the effectiveness in predicting returns, we use two different approaches. First, we calculate the absolute value of the Spearman correlation coefficient between the feature and the monthly returns. We use the absolute value because the correlation is negative for the mean user bot probability (higher bot probability leads to lower returns, as we saw earlier). Second, we try to predict the sign of the returns using the feature. To do this we construct receiver operating characteristic (ROC) curves for each feature and then calculate the corresponding area under the curve (AUC) metric. The AUC is a common metric used to evaluate the predictive power of a binary classifier. In our case, the binary label for each cryptocurrency is the sign of its returns in each month, and the classifier is the feature value. The AUC is one for a perfect classifier and 0.5 for random guessing.

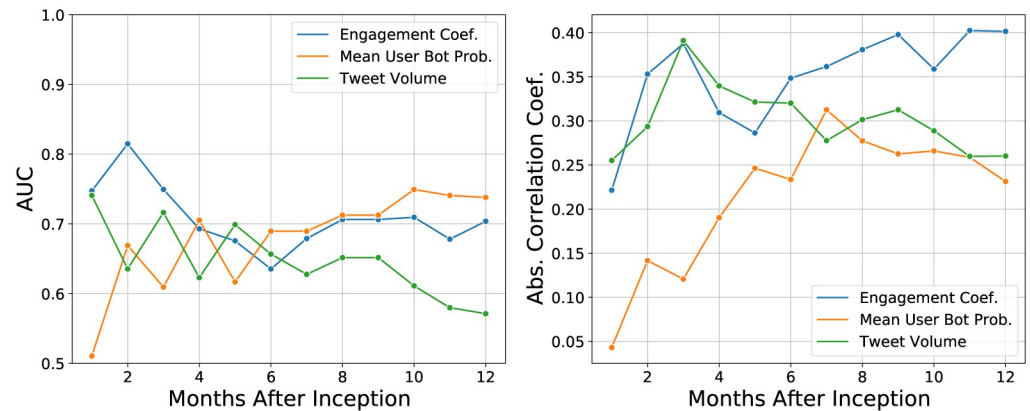


Fig 6. Plot of the (left) AUC and (right) absolute value of the Spearman correlation coefficient for predicting returns versus time (relative to one month after the cryptocurrency's creation) for the social media features.

<https://doi.org/10.1371/journal.pone.0284501.g006>

Fig 6 shows the absolute value of the Spearman correlation coefficient and the AUC versus time for each feature. We see that in the early stages the engagement coefficient and tweet volume have a much higher correlation coefficient and AUC than the bot probability. This suggests that the engagement coefficient and tweet volume are better at predicting initial performance. The AUC and absolute correlation coefficient for the one month returns are nearly equivalent for the engagement coefficient and tweet volume. However, in the later months, the engagement coefficient and mean user bot probability have similar AUC values around 0.7. In contrast, the tweet volume AUCs for returns after six months are below 0.65.

The absolute correlation coefficient of the engagement coefficient and tweet volume are much greater than the mean user bot probability up to six month returns. After six months, tweet volume and mean user bot probability have similar values. However, the engagement coefficient correlation is much greater in this period. For returns up to six months, tweet volume and engagement coefficient have similar correlation coefficients.

Both measures show that the dependence of the one month returns is higher for the engagement coefficient and tweet volume than the mean user bot probability. The tweet volume and engagement coefficient are directly measuring initial social media activity. High enthusiasm for the cryptocurrency on social media early will result in people buying the cryptocurrency, leading to higher early returns. This may be why in the early stages of a cryptocurrency, the enthusiasm of social media users appears to be a better predictor of performance than the presence (or absence) of bots posting content. Of the two social media enthusiasm features, the engagement coefficient seems to have a stronger dependence with the returns than the tweet volume. This suggests that for investment strategies, it may be better to use the engagement coefficient. We explore this in the next section.

Investment strategies

We now consider investment strategies based on social media features. The strategies we consider utilize thresholds for the features. We set a minimum threshold for a feature, and then purchase a fixed dollar amount of any cryptocurrency whose feature value exceeds the threshold. Each cryptocurrency that is purchased is held for a fixed amount of time, after which it is sold. With this approach, the cryptocurrencies are purchased and sold at different times, but held for an equal amount of time. Throughout our analysis, we consider a baseline portfolio

that invests in all of the cryptocurrencies. This portfolio has a feature threshold equal to zero since all features are positive.

In calculating the feature values, we only use data available during the first month of each cryptocurrency's existence, as a real investor would do. This is straightforward for the tweet volume and mean user bot probability. However, to calculate the engagement coefficient, we use the approach where the model is estimated only using prior social media data.

In our first analysis we choose the thresholds to be quartiles of the respective feature distribution for all cryptocurrencies in our dataset. This way, each higher threshold removes a fixed number of cryptocurrencies from the portfolio. We note that this approach could not be used in practice, as the thresholds can depend on data created after the purchase date of a cryptocurrency. Therefore, this analysis does not tell us what feature performs the best in actual portfolios. Rather, we use this approach to understand the properties of the different social media features with respect to the returns.

We construct a portfolio for each feature and quartile threshold. For three social media features and five thresholds (four quartiles plus the baseline portfolio) this gives 15 different portfolios. We can then look at the returns of these portfolios for different holding times. Fig 7 shows the returns versus the cryptocurrency holding time for each portfolio, grouped by the feature. Each curve corresponds to a threshold for the feature.

The returns for each feature exhibit different behaviors with respect to the thresholds. For the tweet volume, the highest post-six month return is achieved for a threshold of 35,800 tweets, which is the 75th quantile. In contrast, the engagement coefficient achieves the highest post-six month return at a threshold of 1.01×10^{-4} , which is the median. The 75th quantile does the worst for returns with holding times greater than two months. The mean user bot probability portfolios all perform equal to or worse than the baseline portfolio. This analysis suggests that the bot probability is not as useful for investment strategies compared to the other features.

To further study the effect of the threshold choice, we plot the returns versus the feature threshold in Fig 8. The tweet volume and bot probability show clusters in their respective feature values. There is no clear pattern for these two features in their returns. As the threshold increases, the returns can increase or decrease. In contrast, the engagement coefficient shows much clearer patterns. The feature values are not clustered, unlike the tweet volume and bot probability. We see that for a one month holding time, the returns increase monotonically in the threshold. For other holding times, the returns show a monotonic increase up to a threshold of 10^{-4} , after which the returns decrease. Cryptocurrencies with engagement coefficients above this threshold appear to have low long-term returns, but high short-term (one month) returns. This supports our earlier observations in Fig 6 that the engagement coefficient may be a useful feature for selecting cryptocurrencies which perform well over short time horizons.

We now try to recreate the scenario faced by an investor and determine the threshold using historical data. We first select an investment date. Then, we use quartiles of the feature distribution for all cryptocurrencies created before this date as threshold values. We use the same investment strategy as before, but now only consider cryptocurrencies created after the investment date. We consider returns for short-term holding times (one month) and long-term holding times (one year).

The resulting returns for different features, thresholds, and investment dates are shown in Figs 9 and 10. For the tweet volume and engagement coefficient, the one month returns are positive for all investment dates and thresholds. In contrast, the bot probability has negative returns for different thresholds and investment dates. The highest returns for each investment date are achieved by the engagement coefficient using the 75th quantile as the threshold. The corresponding threshold values ranged from 2.13×10^{-4} to 2.89×10^{-4} . On the upper end, the

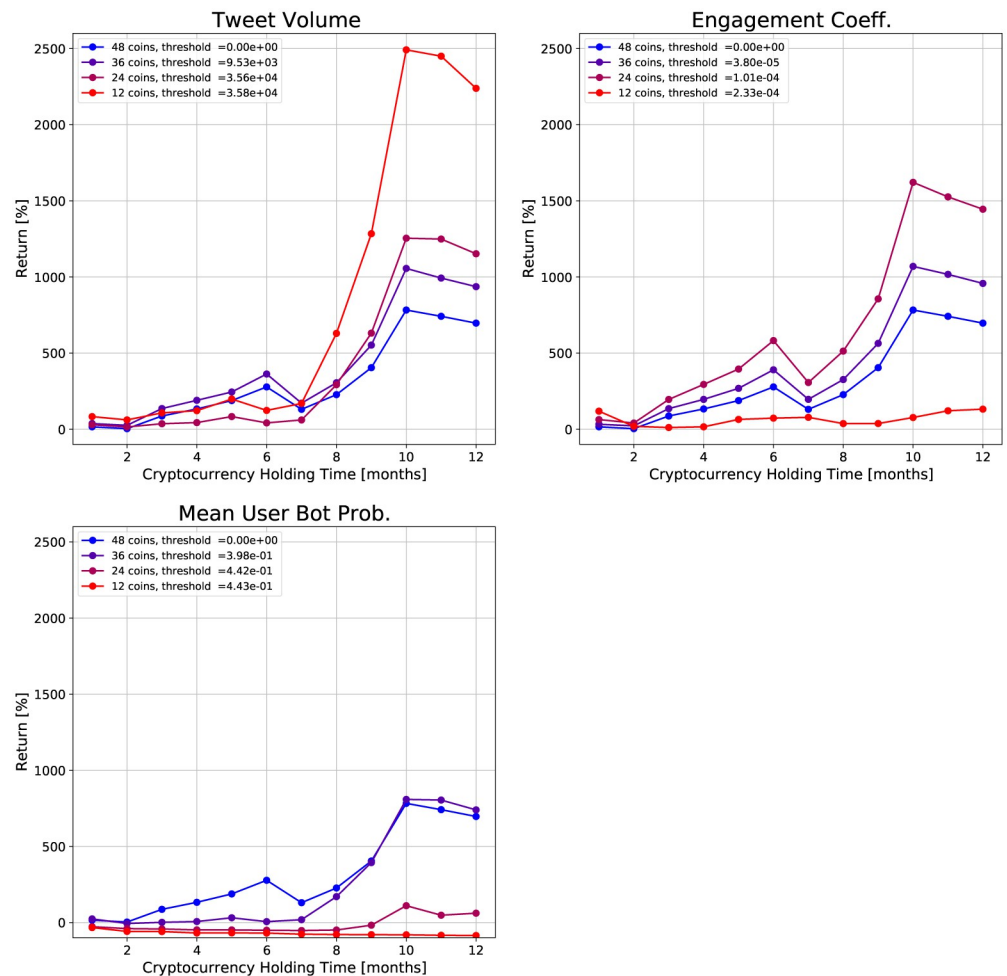


Fig 7. Plots of returns of investment strategies versus holding time for different social media features. The investment strategies are based on fixed cryptocurrency holding times and minimum feature thresholds. Each plot is titled with the feature considered.

<https://doi.org/10.1371/journal.pone.0284501.g007>

returns for the engagement coefficient portfolios are nearly 200%. In contrast, the maximum return for the tweet volume is below 150%, and for the bot probability it below 120%. These results show that the engagement coefficient is the best social media feature for maximizing short-term returns.

The one year returns are negative for investments made in 2021. There is not a substantial difference in returns with respect to the feature threshold when the returns are negative. For earlier investment dates, no feature or threshold shows clear dominance. These findings suggest that long-term returns are less dependent on social media features than short-term returns.

Discussion

The social media activity surrounding certain topics can be predictive of their future performance. Here we found such a relationship for cryptocurrencies. We presented a Poisson model for social media engagement that is simple to estimate, does not require a large amount of data, and can be applied to any topic, not just cryptocurrencies. The model produced a

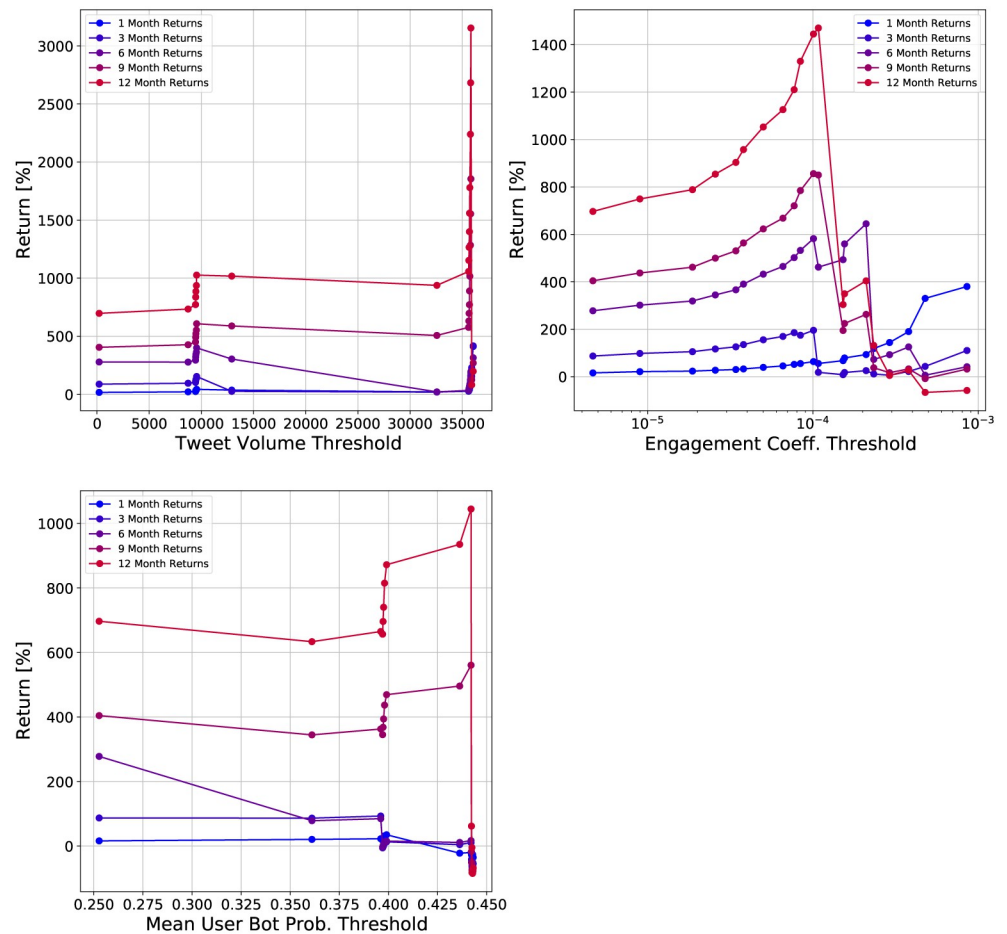


Fig 8. Plots of returns of investment strategies versus feature threshold for different social media features. The investment strategies are based on fixed cryptocurrency holding times and minimum feature thresholds.

<https://doi.org/10.1371/journal.pone.0284501.g008>

numerical value for each cryptocurrency called the engagement coefficient. We found that this value was correlated with the future returns of the cryptocurrencies, and was most predictive of short-term returns. We also found that excessive bot activity surrounding a cryptocurrency was correlated with lower future returns, though the bot activity was less predictive than the engagement coefficient. We developed a simple investment strategy utilizing different social media features, investing only in cryptocurrencies whose feature value exceeded a fixed threshold. We found that if the feature was the engagement coefficient, then this strategy was effective for increasing short-term returns, but not as much for long-term returns. Overall, our results show that one can use social media activity to predict the future performance of cryptocurrencies over longer time horizons than previously known (months versus hours or days) [12].

It is also worth mentioning the potential societal impact of our findings. If it is truly the case that virality significantly drops after 1 month, investors can begin to exit their speculative positions, and potentially avoid the infamous “pump and dump”. As their chosen cryptocurrency begins to lose engagement, their expected returns should be revised downward, which in turn should influence their investment decisions. Given the plethora of investors who have suffered losses to “pump and dump” schemes [38], we believe could have significant societal

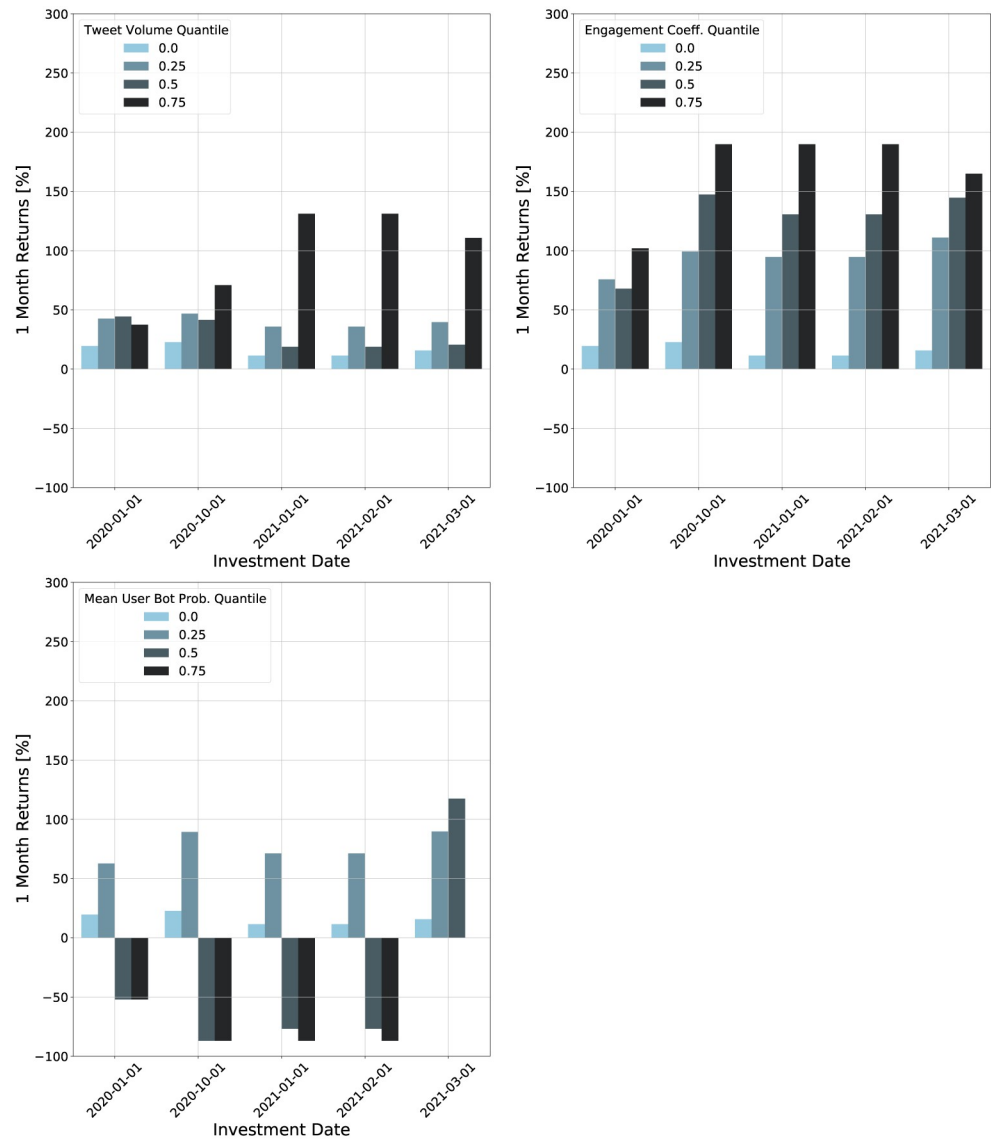


Fig 9. Plots of one month returns of investment strategies for different investment dates based on thresholds for different social media features. The thresholds are quartiles of the relevant feature for cryptocurrencies created before the investment date.

<https://doi.org/10.1371/journal.pone.0284501.g009>

impact. On the flip side, more experienced asset managers could estimate the engagement coefficients using a daily social media feed for a given cryptocurrency, and appropriately make short-term investments in the most “viral” cryptocurrencies.

Limitations

Our work also carries limitations directly associated with cryptocurrency. For example, our study of cryptocurrency ICOs was incidentally in a bull market, in which investors are more likely to be “risk-loving”. Second, although we believe we studied an ample number of coins subject to a minimum fundraising goal, the number of coins available was still limited. This

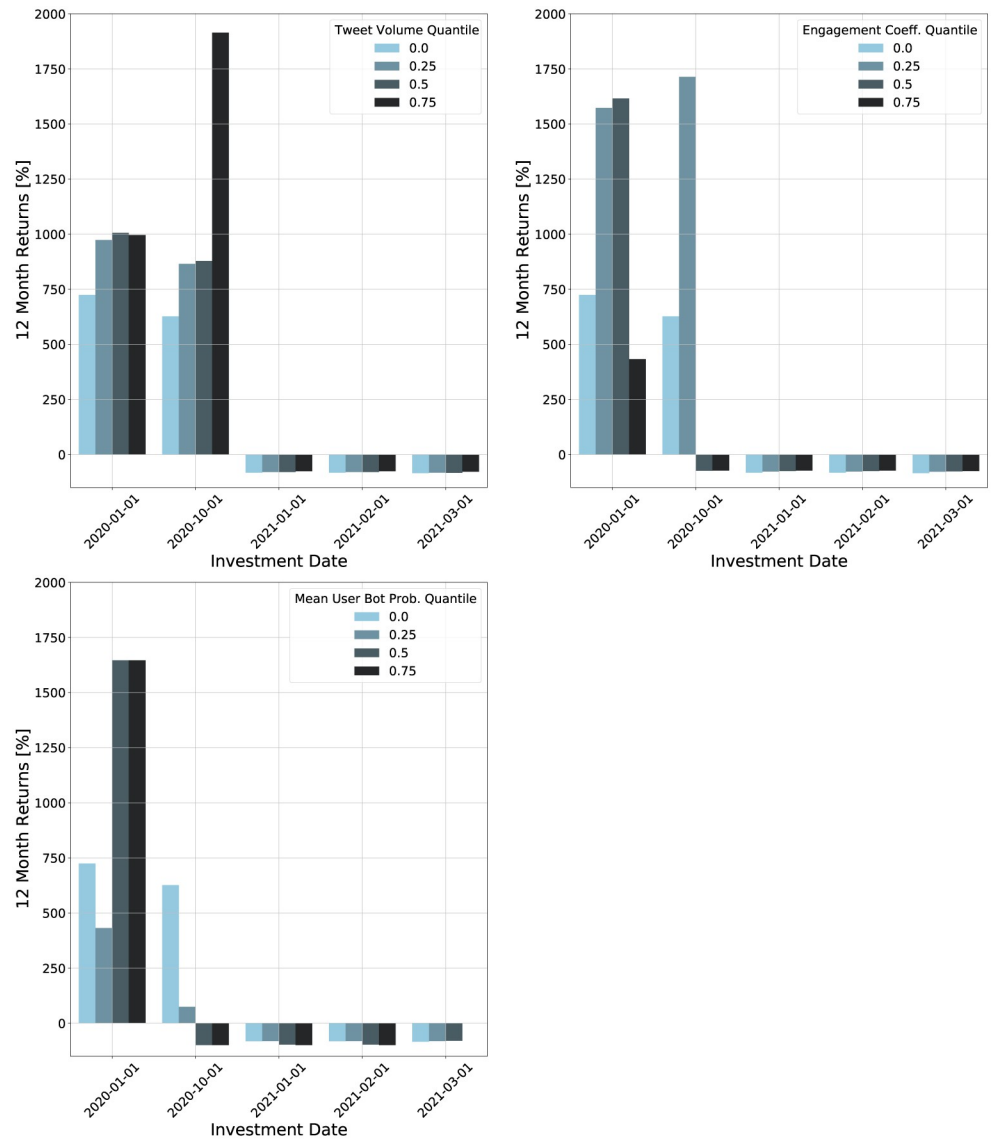


Fig 10. Plots of 12 month returns of investment strategies for different investment dates based on thresholds for different social media features. The thresholds are quartiles of the relevant feature for cryptocurrencies created before the investment date.

<https://doi.org/10.1371/journal.pone.0284501.g010>

was driven mostly by the youth of cryptocurrency as an investment vehicle. We acknowledge these limitations, but believe our work has been exhaustive, conditional on their existence.

Conclusion

The generality of our social media engagement model makes it applicable in diverse settings. Future studies can investigate the predictive power of the model for other topics, such as movies and television shows (ratings, ticket sales), fashion items (sales), political campaigns (votes), or even cryptocurrencies with larger market capitalizations. Indeed, such studies can establish if the range of engagement coefficients we saw here that correlated with positive future returns are similar in other fields and social media platforms. In addition, one can use

our model to estimate time varying engagement coefficients using temporally segmented data to track the evolution of engagement over longer time horizons. The simplicity of our model allows it to be easily applied in this dynamic setting. Finally, we note that our model can be also used to detect artificial manipulation of social media in the context of pump and dump schemes. In this way, our work naturally complements that of [15] and can become another tool to protect investors in the cryptocurrency space.

Supporting information

S1 Data. The data used in this study can be found at <https://github.com/kai-trading-bot/engagement>.

(TXT)

S1 Appendix. Derivation of maximum likelihood estimates.

(PDF)

S1 Table. Cryptocurrency data.

(PDF)

Author Contributions

Conceptualization: Khizar Qureshi.

Data curation: Khizar Qureshi.

Formal analysis: Khizar Qureshi, Tauhid Zaman.

Investigation: Khizar Qureshi, Tauhid Zaman.

Methodology: Khizar Qureshi.

Project administration: Tauhid Zaman.

Resources: Tauhid Zaman.

Software: Khizar Qureshi.

Supervision: Tauhid Zaman.

Validation: Khizar Qureshi, Tauhid Zaman.

Visualization: Khizar Qureshi, Tauhid Zaman.

Writing – original draft: Khizar Qureshi, Tauhid Zaman.

Writing – review & editing: Khizar Qureshi, Tauhid Zaman.

References

1. Howarth J. How Many Cryptocurrencies Are There In 2022?; 2022. <https://explodingtopics.com/blog/number-of-cryptocurrencies>.
2. Coin Market Cap. Top Cryptocurrency Spot Exchanges; 2022. <https://coinmarketcap.com/rankings/exchanges/>.
3. Wüstenfeld J, Geldner T. Economic uncertainty and national bitcoin trading activity. *The North American Journal of Economics and Finance*. 2022; 59:101625. <https://doi.org/10.1016/j.najef.2021.101625>
4. Almeida J, Gonçalves TC. A systematic literature review of investor behavior in the cryptocurrency markets. *Journal of Behavioral and Experimental Finance*. 2023; p. 100785. <https://doi.org/10.1016/j.jbef.2022.100785>
5. Goldman D. What is Dogecoin? How a joke became hotter than bitcoin; 2021. <https://edition.cnn.com/2021/04/17/investing/what-is-dogecoin/index.html/>.

6. Ante L. How Elon Musk's twitter activity moves cryptocurrency markets. Available at SSRN 3778844. 2021;.
7. Hart ea J. Attachment theory as a framework for explaining engagement with Facebook. *Personality and Individual Differences*. 2015;.
8. He ea K. Practical Lessons from Predicting Clicks on Ads at Facebook. *International Workshop on Data Mining for Online Advertising*. 2022;.
9. Luarn ea P. Influence of Facebook brand-page posts on online engagement. *Online Information Review*. 2015;.
10. Malhotra ea A. How to Create Brand Engagement on Facebook. *MIT Sloan Review*. 2013;.
11. Zhang J, Zhang C. Do cryptocurrency markets react to issuer sentiments? Evidence from Twitter. *Research in International Business and Finance*. 2022; 61:101656. <https://doi.org/10.1016/j.ribaf.2022.101656>
12. Steinert L, Herff C. Predicting altcoin returns using social media. *PloS one*. 2018; 13(12):e0208119. <https://doi.org/10.1371/journal.pone.0208119> PMID: 30513110
13. Garcia D, Schweitzer F. Social signals and algorithmic trading of Bitcoin. *Royal Society open science*. 2015; 2(9):150288. <https://doi.org/10.1098/rsos.150288> PMID: 26473051
14. Wu W, Tiwari AK, Gozgor G, Leping H. Does economic policy uncertainty affect cryptocurrency markets? Evidence from Twitter-based uncertainty measures. *Research in International Business and Finance*. 2021; 58:101478
15. Nghiem H, Muric G, Morstatter F, Ferrara E. Detecting cryptocurrency pump-and-dump frauds using market and social signals. *Expert Systems with Applications*. 2021; 182:115284. <https://doi.org/10.1016/j.eswa.2021.115284>
16. Kyriazis ea N. The differential influence of social media sentiment on cryptocurrency returns and volatility during COVID-19. *Quarterly Review of Economics and Finance*. 2022;.
17. Asur S, Huberman BA. Predicting the future with social media. In: 2010 IEEE/WIC/ACM international conference on web intelligence and intelligent agent technology. vol. 1. IEEE; 2010. p. 492–499.
18. Bollen J, Mao H, Zeng X. Twitter mood predicts the stock market. *Journal of computational science*. 2011; 2(1):1–8. <https://doi.org/10.1016/j.jocs.2010.12.007>
19. Zhang X, Fuehres H, Gloor PA. Predicting stock market indicators through twitter "I hope it is not as bad as I fear". *Procedia-Social and Behavioral Sciences*. 2011; 26:55–62. <https://doi.org/10.1016/j.sbspro.2011.10.562>
20. Sprenger TO, Tumasjan A, Sandner PG, Welpe IM. Tweets and trades: The information content of stock microblogs. *European Financial Management*. 2014; 20(5):926–957. <https://doi.org/10.1111/j.1468-036X.2013.12007.x>
21. Mittal A, Goel A. Stock prediction using twitter sentiment analysis. Stanford University, CS229 (2011 <http://cs229.stanford.edu/proj2011/GoelMittal-StockMarketPredictionUsingTwitterSentimentAnalysis.pdf>). 2012;15:2352.
22. Rao T, Srivastava S, et al. Analyzing stock market movements using twitter sentiment analysis. 2012;.
23. Si J, Mukherjee A, Liu B, Li Q, Li H, Deng X. Exploiting topic based twitter sentiment for stock prediction. In: *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*; 2013. p. 24–29.
24. Sul HK, Dennis AR, Yuan L. Trading on twitter: Using social media sentiment to predict stock returns. *Decision Sciences*. 2017; 48(3):454–488. <https://doi.org/10.1111/deci.12229>
25. Bartov E, Faurel L, Mohanram PS. Can Twitter help predict firm-level earnings and stock returns? *The Accounting Review*. 2018; 93(3):25–57. <https://doi.org/10.2308/accr-51865>
26. Oliveira N, Cortez P, Areal N. 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*. 2017; 73:125–144. <https://doi.org/10.1016/j.eswa.2016.12.036>
27. Bandopadhyaya JAL A. Measuring Investor Sentiment in Equity Markets. *Journal of Asset Management*. 2006;.
28. Zaman T, Fox EB, Bradlow ET. A bayesian approach for predicting the popularity of tweets. *The Annals of Applied Statistics*. 2014; 8(3):1583–1611. <https://doi.org/10.1214/14-AOAS741>
29. CoinMarketCap. Today's Cryptocurrency Prices by Market Cap; 2022. <https://coinmarketcap.com/>.
30. CoinGecko. Cryptocurrency Prices by Market Cap; 2022. <https://www.coingecko.com/>.
31. Twitter. Academic Research access; 2022. <https://developer.twitter.com/en/products/twitter-api/academic-research>.
32. Zirlin J. Axie Infinity; 2018 (accessed 2022). <https://axieinfinity.com/>.

33. Varol O, Ferrara E, Davis C, Menczer F, Flammini A. Online human-bot interactions: Detection, estimation, and characterization. In: Proceedings of the international AAAI conference on web and social media. vol. 11; 2017. p. 280–289.
34. Yang KC, Ferrara E, Menczer F. Botometer 101: Social bot practicum for computational social scientists. arXiv preprint arXiv:220101608. 2022;.
35. Vosoughi S, Roy D, Aral S. The spread of true and false news online. *science*. 2018; 359(6380):1146–1151. <https://doi.org/10.1126/science.aap9559> PMID: 29590045
36. Torusdağ MB, Kutlu M, Selçuk AA. Are we secure from bots? Investigating vulnerabilities of botometer. In: 2020 5th International Conference on Computer Science and Engineering (UBMK). IEEE; 2020. p. 343–348.
37. Menczer ea F. Botometer: An OSoMe project; 2021 (accessed 2022). <http://www.https://botometer.osome.iu.edu/>.
38. Kamps KB J. To the moon: defining and detecting cryptocurrency pump-and-dumps. *Crime Science*. 2018;.