Cashtag Piggybacking: Uncovering Spam and Bot Activity in Stock Microblogs on Twitter

STEFANO CRESCI, Institute of Informatics and Telematics, IIT-CNR, Italy
FABRIZIO LILLO, Department of Mathematics, University of Bologna, Italy and Scuola Normale Superiore of Pisa, Italy
DANIELE REGOLI, Azimut Analytics srl, Milano, Italy and Scuola Normale Superiore of Pisa, Italy
SERENA TARDELLI and MAURIZIO TESCONI, Institute of Informatics and Telematics, IIT-CNR, Italy

Microblogs are increasingly exploited for predicting prices and traded volumes of stocks in financial markets. However, it has been demonstrated that much of the content shared in microblogging platforms is created and publicized by bots and spammers. Yet, the presence (or lack thereof) and the impact of fake stock microblogs has never been systematically investigated before. Here, we study 9M tweets related to stocks of the five main financial markets in the US. By comparing tweets with financial data from Google Finance, we highlight important characteristics of Twitter stock microblogs. More importantly, we uncover a malicious practice—referred to as *cashtag piggybacking*—perpetrated by coordinated groups of bots and likely aimed at promoting low-value stocks by exploiting the popularity of high-value ones. Among the findings of our study is that as much as 71% of the authors of suspicious financial tweets are classified as bots by a state-of-the-art spambot-detection algorithm. Furthermore, 37% of them were suspended by Twitter a few months after our investigation. Our results call for the adoption of spam- and bot-detection techniques in all studies and applications that exploit user-generated content for predicting the stock market.

CCS Concepts: • **Information systems** → **Social networks**; • **Security and privacy** → **Social network security and privacy**; • **Applied computing** → Economics;

Additional Key Words and Phrases: Social spam, social networks security, spam and bot detection, stock market, Twitter

ACM Reference format:
Stefano Cresci, Fabrizio Lillo, Daniele Regoli, Serena Tardelli, and Maurizio Tesconi. 2019. Cashtag Piggybacking: Uncovering Spam and Bot Activity in Stock Microblogs on Twitter. *ACM Trans. Web* 13, 2, Article 11 (April 2019), 27 pages.
https://doi.org/10.1145/3313184

This research is supported in part by the EU H2020 Program INFRAIA-1-2014-2015: Research Infrastructures under Grant No.: 654024 *SoBigData: Social Mining & Big Data Ecosystem.*
Authors' addresses: S. Cresci (corresponding author), S. Tardelli, and M. Tesconi, Via G. Moruzzi 1, 56124, Pisa, Italy; emails: {stefano.cresci, serena.tardelli, maurizio.tesconi}@iit.cnr.it; F. Lillo, Via Zamboni 33, 40126 Bologna, Italy; email: fabrizio.lillo@unibo.it; D. Regoli, Azimut Analytics srl, Foro Buonaparte 24, 20121, Milano, Italy; email: daniele.regoli@sns.it.

1559-1131/2019/04-ART11 $15.00
https://doi.org/10.1145/3313184

1 INTRODUCTION

The exploitation of user-generated content in microblogs for the prediction of real-world phenomena has recently gained huge momentum [24, 66]. An important application domain is that of finance and, in particular, stock-market prediction. Indeed, a number of works developed algorithms and tools for extracting valuable information (e.g., sentiment scores) from microblogs and proved capable of predicting prices and traded volumes of stocks in financial markets [12]. Notably, finance is increasingly relying on this information through the development of automatic trading systems.

All such works are grounded on the assumption that microblogs collectively represent a reliable proxy for the opinions of masses of users. Meanwhile, however, evidence of fake accounts as well as spam and automated (bot) activities in social platforms is being reported at a growing rate [25, 37]. The existence of fictitious, synthetic content appears to be pervasive, since it has been witnessed both in online discussions about important societal topics (e.g., politics, terrorism, immigration, health-care), as well as in discussions about seemingly less relevant topics, such as products on sale on e-commerce platforms and mobile applications [28]. For instance, regarding health-related topics, it has been demonstrated that bots are exploited to promote online health content to legitimize the vaccine debate and disseminate anti-vaccine messages [3, 14]. In addition, there has been evidence of social bots encouraging electronic cigarette consumption through the use of false first-person experiences and feedback [22, 73]. Regarding politics, it has been demonstrated that bots tampered with recent US [9], Italian [26], French [35], Japanese [65], and—to a minor extent—German [13, 52] political elections, as well as with online discussions about the 2016 UK Brexit referendum [8].

Thus, on the one hand, user-generated content in microblogs is being exploited for predicting trends in the stock market; on the other hand, without a thorough investigation, we run the risk that much of the content we rely on is actually fake and possibly purposely created to mislead algorithms and users alike [60]. Should this risk materialize, real-world consequences would be severe, as already anticipated by a few noteworthy events [34]. On May 6, 2010, the Dow Jones Industrial Average had the biggest one-day drop in history, later called the *Flash Crash*. After five months, an investigation concluded that one of the possible causes was an automated high-frequency trading system that had incorrectly assessed some information collected from the web [45]. In 2013, the US International Press Officer's Twitter account was hacked and a false rumor was posted reporting that President Obama was injured during a terrorist attack. The fake news rapidly caused a stock-market collapse that burned $136B.[1] Then, in 2014, the unknown *Cynk Technology* briefly became a $6B company: Automatic trading algorithms detected a fake social discussion and began to invest heavily in the company's shares. By the time analysts noticed the orchestration, investments had already turned into heavy losses.[2]

1.1 Contributions

In a recent investigation [30], we reported the first preliminary evidence of the presence of financial spam in stock microblogs, raising serious concerns over the reliability of such information. Here, we deepen our previous analyses by performing a number of additional experiments on co-occurring cashtags, on financial markets, and on suspicious users. Specifically, we extend our previous work with the following novel and unpublished contributions:

- We analyze co-occurring cashtags in financial tweets by focusing on their industrial and economic classification. In detail, we show that co-occurrences of stocks in suspicious

[1] http://www.telegraph.co.uk/finance/markets/10013768/Bogus-AP-tweet-about-explosion-at-the-White-House-wipes-billions-off-US-markets.html.
[2] http://mashable.com/2014/07/10/cynk/#HD9o6llp6gqw.

tweets are not motivated by the fact that those stocks belong to the same industrial or economic sectors (§ 5.2);
- Since real-world relatedness (as expressed by industrial classification) is not a plausible explanation for co-occurring stocks, we then turn our attention to market capitalization. We demonstrate that, in suspicious tweets, high-capitalization companies co-occur with low-capitalization ones. Moreover, we show that this large difference can not be explained by the intrinsic characteristics of our dataset, but it is rather the consequence of an external action (§ 5.3);
- We compare the social and financial importance of investigated companies, highlighting that stocks of one specific market (OTCMKTS) feature a suspiciously high social importance despite their low financial importance. This result is in contrast to measurements obtained for stocks of the other markets—e.g., NASDAQ, NYSE, NYSEARCA, and NYSEMKT (§ 5.4);
- We employ a state-of-the-art spambot-detection technique to analyze authors of suspicious tweets. Results show that 71% of suspicious users are classified as bots. Furthermore, 37% of them were also suspended by Twitter a few months after our investigation (§ 6).

Summarizing, this study moves in the direction of investigating the presence of spam and bot activity in stock microblogs, thus paving the way for the development of intelligent financial-spam filtering techniques. To reach our goal, we first collect a rich dataset of 9M tweets posted between May and September 2017, discussing stocks of the five main financial markets in the US. We enrich our dataset by collecting financial information from Google Finance about the 30,032 companies mentioned in our tweets. Cross-checking discussion patterns on Twitter against official data from Google Finance uncovers anomalies in tweets related to some low-value companies. Further investigation of this issue reveals a large-scale speculative campaign—which we refer to as *cashtag piggybacking*—perpetrated by coordinated groups of bots and aimed at promoting low-value stocks by exploiting the popularity of high-value ones. Finally, we analyze a subset of authors of suspicious tweets with state-of-the-art bot-detection techniques, identifying 71% (18,509 accounts) of them as bots.

1.2 Cashtag Piggybacking

Results of our study uncover a large presence of bot accounts in stock microblogs on Twitter. More specifically, we thoroughly document a practice aimed at promoting low-cap stocks (mainly OTCMKTS stocks) by exploiting the popularity of high-cap ones.

We name this novel kind of spam as *cashtag piggybacking*, by borrowing the concept of *piggyback*[3] from the field of computer networks [70]. In many network protocols, a sender node is allowed to deliver short messages (e.g., ACKs to previous packets) to a receiver node, without sending a dedicated packet. In fact, the sender can postpone the short message until a new packet must be sent. At this time, the sender piggybacks (i.e., adds) the message as part of the outgoing packet. In network protocols, piggybacking allows to increase the efficiency in communications [70]. Indeed, fewer packets are sent, since small amounts of information can be sent "on top of the shoulders" of other packets.

Within the context of stock microblogs, we show that coordinated groups of bots piggyback some low-value stocks "on top of the shoulders" of other high-value stocks. Hence, the *cashtag piggybacking* name.

[3] https://en.wikipedia.org/wiki/Piggybacking_(data_transmission).

1.3 Roadmap

The remainder of this article is organized as follows: Section 2 discusses relevant related work in stock-market prediction from social media and in spam and bot detection. Then, Section 3 describes the dataset used in this study. In Section 4, we briefly provide an overview of the characteristics of our dataset, and we describe the methodology adopted to identify suspicious tweets. In Section 5, we analyze suspicious tweets and financial markets from several viewpoints. Instead, in Section 6, we turn our attention to the authors of the suspicious tweets, looking for bots among them. Section 7 provides a critical discussion of our results. Finally, Section 8 draws conclusions and highlights some promising directions for future research and experimentation.

2 RELATED WORK

Since no study has previously addressed bot activity in stock microblogs, this section is organized to separately survey previous work either related to the exploitation of user-generated content for financial purposes or to spam and bot characterization.

2.1 Finance and Social Media

Works in this field are based on the idea underlying the Hong-Page theorem [44]. Such a theorem, when cast in the financial domain, states that user-generated messages about a company's future prospects provide a rich and diverse source of information, in contrast to what the small number of traditional financial analysts can offer.

Starting from the general assumption of the Hong-Page theorem, much effort has been devoted towards the detection of correlations between metrics extracted from social media posts and stock-market prices. In particular, *sentiment* metrics have been widely used as a predictor for stock prices and other economic indicators [11, 21, 38, 42, 62, 69]. The primary role played by the sentiment of the users as a financial predictor is also testified by the interest in developing domain-specific sentiment classifiers for the financial domain [23, 68]. For example, the classifier developed in [15] is based solely on the sentiment analysis of tweets and accurately predicts the next day trend of the stock values of specific companies. Similarly, the one developed in [61] predicts opening and closing stock-market prices with high accuracy. Another study aims to create and improve domain-specific sentiment lexicon and sentiment-oriented word embedding models to help sentiment analysis in the financial domain [55].

Other studies have instead proposed to exploit the overall volume of tweets about a company [59] and the topology of stock networks [64] as predictors of financial performance. Specifically, authors of Reference [59] envisioned the possibility to automatically buy or sell stocks based on the presence of a peak in the volume of tweets. However, subsequent work [83] evaluated the informativeness of sentiment- and volume-derived predictors, showing that the sentiment of tweets contains significantly more information for predicting stock prices than just their volume. The role of *influencers* in social media has also been identified as a strong contributing factor to the formation of market trends [16]. Others have instead used weblogs for studying the relationships between different companies [51]. In detail, co-occurrences of stock mentions in weblogs have been exploited to create a graph of companies, which was subsequently clustered. Authors have verified that companies belonging to the same clusters feature strong correlations in their stock prices. This methodology can be employed for market prediction and as a portfolio-selection method, which has been shown to outperform traditional strategies based on company sectors or historical stock prices.

Another line of research focused on the exploitation of social media content for monitoring and predicting firm equity value. As an example, the study in Reference [81] investigated the effect

of social media and conventional media, their relative importance, and their inter-relatedness on short-term firm equity value prediction. Findings indicated that social media have a stronger relationship with firm equity value than conventional media, while social and conventional media have a strong interaction effect on stock performance. Similarly, in Reference [58], authors focused on the effects of social media–derived metrics compared with conventional online company behavioral metrics. Results derived from autoregressive models suggested that social media–derived metrics (e.g., weblogs and consumer ratings) are significant leading indicators of firm equity value. Even more interestingly, conventional online behavioral metrics (Google searches and web traffic) have a significant yet substantially weaker predictive relationship with firm equity value than social media metrics. Another study [57] from the same authors assessed the extent to which "consumer buzz," in the form of user-generated reviews, recommendations, and blog posts, influence firm value. Results support the dynamic relationships of buzz and web traffic with firm value, and the related mediation effects of buzz and traffic. The study also uncovered significant market competition effects, including effects of both a firm's own and its rivals' buzz and traffic.

Nowadays, results of studies such as those briefly surveyed in this section are leveraged for the development of automatic trading systems that are largely fed with social media–derived information [33]. As a consequence, such automatic systems can potentially suffer severe problems caused by large quantities of fictitious posts. As discussed in the next section, the presence of social bots—and of the fake content they produce—is so widespread as to represent a serious, tangible threat to these, and other, systems [41].

2.2 Spam and Bots in Social Media

Since our study is aimed at verifying the presence and the impact of spam and bot activity in stock microblogs, in this section, we focus on discussing previous work about the characterization and detection of spam and bots in social media.

Many developers of spammer accounts make use of bots to simultaneously and continuously post a great deal of spam content. This is one of the reasons why, despite bots being in rather small numbers when compared to legitimate users, they nonetheless have a profound impact on content popularity and activity in social media [1, 41]. In addition, bots are driven to act in a *coordinated* and *synchronized* way, thus amplifying their effects [63, 82]. Another problem with bots is that they *evolve* over time to evade established detection techniques [32, 72, 78]. Hence, newer bots often feature advanced characteristics that make them way harder to detect with respect to older ones. Recently, a general-purpose overview of the landscape of automated accounts was presented in Reference [36]. This work testifies to the emergence of a new wave of social bots, capable of mimicking human behavior and interaction patterns in social media better than ever before. A subsequent study [28] compared "traditional" and "evolved" bots in Twitter, and demonstrated that the latter are almost completely undetected by platform administrators. Moreover, a crowdsourcing campaign showed that even tech-savvy users are incapable of accurately identifying the evolved bots.

Since bots and spammers evolved, putting in place complex mechanisms to evade existing detection systems [27], scholars and platform administrators tried to keep pace by proposing powerful techniques based on profile [5, 25, 54], posting [7, 53, 77], and network [2, 39, 47, 79] characteristics of the accounts. The study presented in Reference [28], however, demonstrated that also the majority of these bot-detection techniques, which are based on off-the-shelf machine learning algorithms applied for analyzing one account at a time, are unable to effectively detect the evolved bots. To overcome this limitation, a recent stream of research proposed ad hoc detection techniques for the collective analysis of groups of accounts, rather than single accounts [18–20, 26, 29, 46, 48, 49, 80]. These techniques achieved better detection results than previous ones [28] and represent nowadays the last bulwark against pervasive malicious accounts in social media.

Fig. 1. Sample tweets with the $AAPL, $WMT, and $AMZN cashtags.

However, the battle is far from over. Indeed, given this worrying picture, it is not surprising that bots have recently proven capable of influencing the public opinion for many crucial topics [8, 9, 35] and in many different ways, such as by spreading fake news [67] or by artificially inflating the popularity of certain posts [10] and public characters [25]. The combination of automatic trading systems feeding on social media data and the pervasive presence of spam and bots motivates our investigation on the presence of spam and bots in stock microblogs. Moreover, the financial domain has already been proven to have peculiar characteristics with respect to many information-processing tasks (e.g., ranking [17] and filtering [71] content, expert finding [76], etc.) to require ad hoc analyses, such as the one carried out in this work.

3 DATASET

Our dataset for this study is composed of: (i) stock microblogs collected from Twitter, and (ii) financial information collected from Google Finance.

3.1 Twitter Data Collection

Twitter users follow the convention of tagging stock microblogs with so-called *cashtags*. The cashtag of a company is composed of a dollar sign followed by its ticker symbol (e.g., $AAPL is the cashtag of *Apple, Inc.*). Figure 1 shows two sample tweets with the $AAPL, $WMT, and $AMZN cashtags. Similarly to hashtags, cashtags visually highlighted on Twitter's interface can be used as an efficient mean to filter content and to collect data about given companies [43]. For this reason, we based our Twitter data collection on an official list of cashtags. Specifically, we first downloaded a list of 6,689 stocks traded on the most important US markets (e.g., NASDAQ, NYSE) from the official NASDAQ website.[4] Then, we collected all tweets shared between May and September 2017, containing at least one cashtag from the list. Data collection from Twitter has been carried out by exploiting Twitter's Streaming APIs[5] [31]. After collecting five months' worth of data, we ended up with ~9M tweets (of which 22% are retweets), posted by ~2.5M distinct users, as shown in Table 1.

As a consequence of our data collection strategy, every tweet in our dataset contains at least one cashtag from the starting list. However, many collected tweets contain more than one cashtag, many of which are related to companies not included in our starting list. Indeed, overall, we collected data of about 30,032 companies traded across 5 different markets.

3.2 Financial Data Collection

We enriched our Twitter dataset by collecting financial information about each of the 30,032 companies found in our tweets. Financial information has been collected from public company data

[4] http://www.nasdaq.com/screening/company-list.aspx.
[5] https://developer.twitter.com/en/docs/tweets/filter-realtime/overview.

Table 1. Financial and Social Dataset Composition

	financial data			twitter data		
markets	*companies*	*median cap. ($)*	*total cap. ($B)*	*users*	*tweets*	*retweets (%)*
NASDAQ	3,013	365,780,000	10,521	252,587	4,017,158	1,017,138 (25%)
NYSE	2,997	1,810,000,000	28,692	265,618	4,410,201	923,123 (21%)
NYSEARCA	726	245,375,000	2,227	56,101	298,445	157,101 (53%)
NYSEMKT	340	78,705,000	256	22,614	196,545	63,944 (33%)
OTCMKTS	22,956	31,480,000	45,457	64,628	584,169	446,293 (76%)
total	30,032	–	87,152	467,241	7,855,518	1,802,705 (23%)

Total values of *users*, *tweets*, and *retweets* only count distinct items and thus do not equal the sum of previous rows.

Fig. 2. Thomson Reuters Business Classification (TRBC) hierarchical schema.

hosted on Google Finance.[6] Among collected financial information is the *market capitalization* (market cap) of a company and its *industrial classification*.

The capitalization is the total dollar market value of a company. For a given company i, it is computed as the share price $P(s_i)$ times the number of outstanding shares $|s_i|$: $C_i = P(s_i) \times |s_i|$. In our study, we take the market cap of a company into account, since it allows us to compare the financial value of that company with its social media popularity and engagement.[7] In Table 1, we report the median capitalization of the companies for each considered market. As shown, important markets such as NYSE and NASDAQ trade, on average, stocks with higher capitalization than those traded in minor markets.

Industrial classification is expressed via the Thomson Reuters Business Classification[8] (TRBC). As shown in Figure 2, TRBC is a 5-level hierarchical sector and industry classification, widely used in the financial domain for computing sector-specific indices. At the topmost (coarse-grained) level, TRBC classifies companies into 10 economic sectors, while at the lowest (fine-grained) level companies are divided into 837 different activities. A few examples of the TRBC industrial classification are reported in Table 2. In our study, we compare companies belonging to the same category across all 5 levels of TRBC.

4 ANALYSIS OF STOCK MICROBLOGS

4.1 Dataset Overview

Surprisingly, the vast majority (76%) of companies mentioned in our dataset do not belong to the NASDAQ list and are traded in OTCMKTS, as shown in Table 1. Having so many OTCMKTS companies

[6] https://www.google.com/finance.
[7] In the remainder, share prices and market capitalizations are considered as of July 4, 2017.
[8] https://financial.thomsonreuters.com/en/products/data-analytics/market-data/indices/trbc-indices.html.

Table 2. Examples of TRBC Classifications

		TRBC levels				
ticker	**company**	*activity*	*industry*	*industrial group*	*business sector*	*economic sector*
AAPL	Apple, Inc.	Computer Hardware-NEC	Computer Hardware	Computers, Phones & Household Electronics	Technology Equipment	Technology
GOOG	Alphabet, Inc.	Search Engines	Internet Services	Software & IT Services	Software & IT Services	Technology
JNJ	Johnson & Johnson	Pharmaceutic-NEC	Pharmaceutic	Pharmaceutic	Pharmaceutics & Medical Research	Healthcare

in our dataset is already an interesting finding, considering that our data collection grounded on a list of high-capitalization (high-cap) companies. OTCMKTS is a US financial market for over-the-counter transactions, and, thus, it has far less stringent requirements than those needed from NASDAQ, NYSE, NYSEARCA, and NYSEMKT. For this reason, many small companies opt to be traded in OTCMKTS instead of the more rigid markets. However, in addition to small-cap companies, OTCMKTS also trades *American depositary receipts* (ADRs),[9] which allow non-US companies to trade stocks in US markets. Otherwise, non-US companies would only be traded in other foreign markets (e.g., stocks of *Samsung Electronics Co., Ltd.* would only be traded in the Korea Exchange). OTCMKTS also trades *Convertible Preferred stocks*,[10] which are a particular kind of stock that give more guarantees to investors with respect to common stocks. Other types of assets might be traded in this market. In our study, we do not discriminate between different types of assets traded in OTCMKTS, and we rely on the financial information contained in Google Finance, irrespective of the kind of traded stocks.

Thus, from a company viewpoint, our dataset is dominated by stocks traded in OTCMKTS. However, OTCMKTS companies play a marginal role from both a financial and social viewpoint, having low median capitalization and small numbers of tweets, the vast majority of which are retweets. In contrast, companies from NASDAQ and NYSE have high capitalization and are mentioned in many tweets, with a low percentage of retweets.

In the following, we report on some of the general characteristics of our dataset. Figure 3(a) shows a cashtag-cloud representing the most tweeted companies in our dataset. In Figure 3(a), cashtags are color-coded to visually highlight companies traded in different markets. The most-tweeted companies in our dataset are in line with recent trends (e.g., the $BTC (*Bitcoin*) and $ETH (*Ethereum*) cryptocurrencies) and with findings of previous works [4, 43] (e.g., $AAPL leading the way, followed by $AMZN, $FB, and $TSLA). Notably, no company from OTCMKTS appears among top-mentioned companies, but instead they play a rather marginal role. Figure 3(b) shows the mean volume of tweets collected per hour. The largest surge of tweets occurs between 10 a.m. and 5 p.m. (US Eastern time), which almost completely overlaps with the opening hours of the New York Stock Exchange (9:30 a.m. to 4 p.m.). This fact further highlights the strong relation between stock microblogs and the real-world stock market. Finally, as previously introduced, many stock microblogs contain more than one cashtag (e.g., the right-hand side tweet in Figure 1). Figure 3(c) shows the distribution of distinct cashtags per tweet, with a mean value of 2 cashtags/tweet. To perform a more detailed analysis, we also compared the mean volume of tweets collected per hour and day of week, with the mean number of cashtags per tweet, by hour and day of the week. Figures 3(d) and 3(e) show them, respectively. In Figure 3(d), we observe that most of the finan-

[9] https://en.wikipedia.org/wiki/American_depositary_receipt.
[10] https://en.wikipedia.org/wiki/Preferred_stock.

(a) Cashtag-cloud of most tweeted companies.

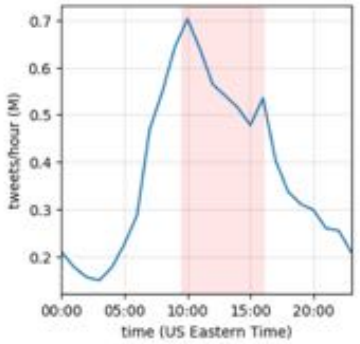

(b) Mean tweet volume per hour. Peak hours overlap with the opening hours of the New York Stock Exchange (red band).

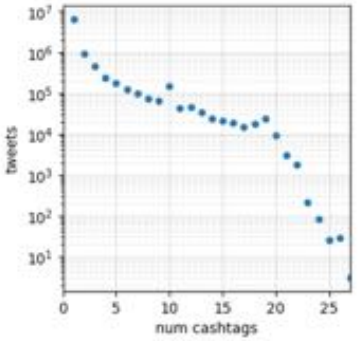

(c) Distribution of the number of cashtags per tweet.

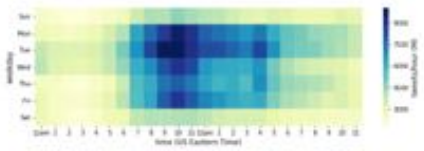

(d) Mean tweet volume per hour and weekday.

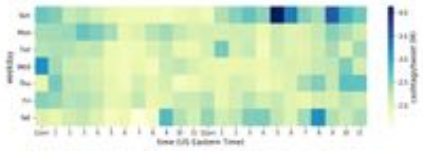

(e) Mean cashtags/tweet per hour and weekday.

Fig. 3. Overall statistics about our dataset.

cial discussions on Twitter take place during working days, around the market opening hours (9:30 a.m.). The same pattern can be seen in Figure 3(b). Figure 3(e) shows that the mean number of cashtags per tweet is very low, reflecting the findings in Figure 3(c), where the mean value is 2 cashtags/tweet. Moreover, it shows a slight change in the number of cashtags per tweet, depending on the time. Indeed, before and after the market hours, it appears that the mean number of cashtags is slightly greater than the one during the market hours.

4.2 Stock Time Series Analysis

To uncover possible malicious behaviors related to stock microblogs, we carry out a fine-grained analysis of our data. Specifically, we build and analyze the hourly time series of each of the 6,689

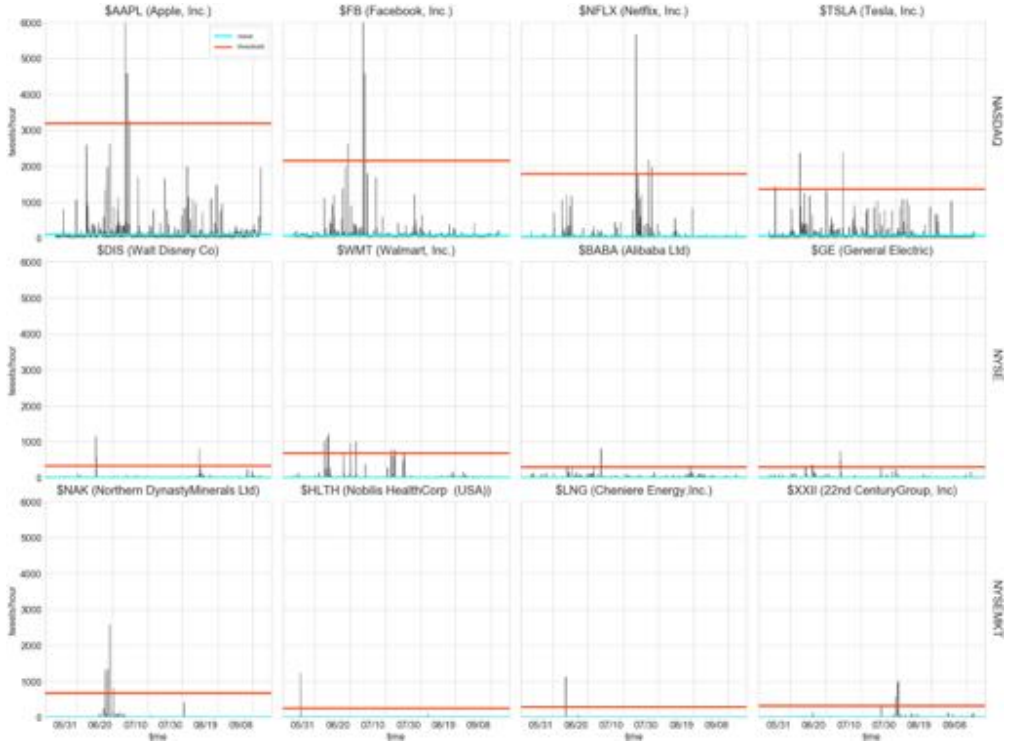

Fig. 4. Examples of stock time series for 12 highly tweeted stocks. Mean values are marked with cyan solid lines and thresholds above which peaks are detected ($K = 10$) are marked with red solid lines.

stocks downloaded from the NASDAQ website. Given a stock i, its time series is defined as $\mathbf{s}_i = (s_{i,1}, s_{i,2}, \ldots, s_{i,N})$, with $s_{i,j}$ being the number of tweets that mentioned the stock i during the hour j. Figure 4 shows some examples of our stock time series for 12 highly tweeted stocks across 3 markets (NASDAQ, NYSE, and NYSEMKT). As shown in Figure (4), stock time series are characterized by long time spans over which tweet discussion volumes remain rather low, occasionally interspersed by large discussion spikes. This behavior is consistent with what has been previously observed in Twitter for other phenomena (e.g., communication patterns related to emergency events [6]). Indeed, the *bursty* and *spiky* characteristics of social communications have been recently explained as a direct consequence of human dynamics [50].

To give a better characterization of this phenomenon, we ran a simple anomaly-detection technique on all the 6,689 time series. As typically done in many time series analysis tasks, our anomaly-detection technique is designed to detect a peak $p_{i,j}$ in a time series $\mathbf{s}_i$ *iff* the tweet volume for the hour j deviates from the mean tweet volume $\bar{\mathbf{s}}_i$ by a number K of standard deviations:

$$p_{i,j} \iff s_{i,j} > \bar{\mathbf{s}}_i + K \times \sigma(\mathbf{s}_i).$$

The parameter K determines the number of peaks found by our anomaly-detection technique. In fact, a bigger K implies that a larger deviation from the mean is needed to detect a peak. Figure 5 shows the number of peaks detected in our time series as a function of the parameter K. For the remainder of our analysis, we set $K = 10$, which represents a trade-off between the height of considered peaks and the number of peaks to analyze. This choice of K results in 1,926 peaks

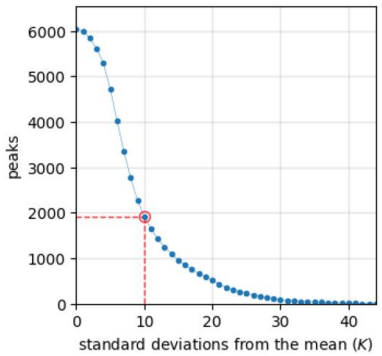

Fig. 5. Number of peaks detected, as a function of K.

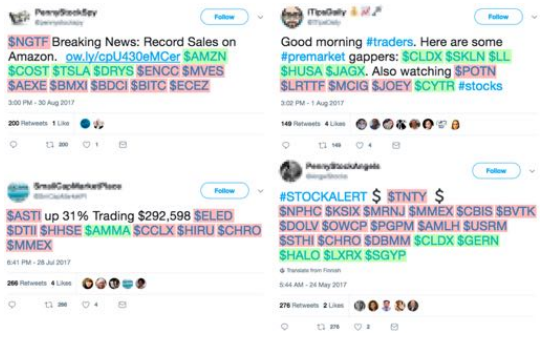

Fig. 6. Examples of suspicious peak tweets. In every tweet, a few cashtags of high-cap stocks (green-colored) co-occur with many cashtags of low-cap stocks (red-colored).

detected in our time series. Time series depicted in Figure 4 also show mean values (cyan solid line) and the 10σ threshold (red solid line) above which peaks are detected.

Next, we are interested in analyzing the tweets that generated the peaks (henceforth, *peak tweets*). In detail, a peak $p_{i,j}$ is composed of a set of tweets $\mathbf{t}_{i,j}$, such that each tweet $t \in \mathbf{t}_{i,j}$ contains the cashtag related to the stock i and has been posted during the hour j (i.e., the peak hour):

$$\mathbf{t}_{i,j} = \left\{ t^1_{i,j}, t^2_{i,j}, \ldots, t^M_{i,j} \right\}, \quad M = s_{i,j}.$$

Thus, for each of the 1,926 peaks $p_{i,j}$, we analyze the corresponding set of tweets $\mathbf{t}_{i,j}$. We find out that, on average, 60% of tweets $t \in \mathbf{t}$ are retweets. In other words, the peaks identified by our anomaly-detection technique are largely composed of retweets. In addition, considering that our time series have hourly granularity, those retweets also occurred within a rather limited time span, in a *bursty* fashion. This finding is particularly interesting also considering that in all our dataset, we had only 23% retweets, versus 60% measured for peak tweets.

We also analyzed tweets $t \in \mathbf{t}$ by considering the co-occurrences of stocks. From this analysis, we see that tweets $t \in \mathbf{t}$ typically contain many more cashtags than tweets $t \notin \mathbf{t}$. The cashtags that co-occur in peak tweets seem unrelated, and the authors of those tweets don't provide further information to explain such co-occurrences. As an example, Figure 6 shows four such suspicious tweets: in every tweet, a few cashtags of high-capitalization (high-cap) stocks co-occur with many cashtags of low-cap stocks. The distributions of the number of retweets per tweet, and of the number of cashtags per tweet, are shown in Figures 7(a) and 7(b), respectively. The distributions are shown with beanplots and allow comparison of values measured for the whole dataset (green-colored) with those measured only in peak tweets (light-blue-colored).

The characteristics of peak tweets previously highlighted—that is, the percentage of retweets and the number of co-occurring cashtags—differ significantly from those measured for the whole dataset. The reason for this peculiar phenomenon could be related to some real-world news or event that motivates the surge of retweets and the co-occurrences of different cashtags. However, such differences could also be the consequence of shady, malicious activity. Indeed, there have already been reports of large groups of bots that coordinately and simultaneously alter popularity and engagement metrics of Twitter users and content [10, 37]. In particular, mass retweets have been identified as one mean to artificially increase the popularity of certain content [28].

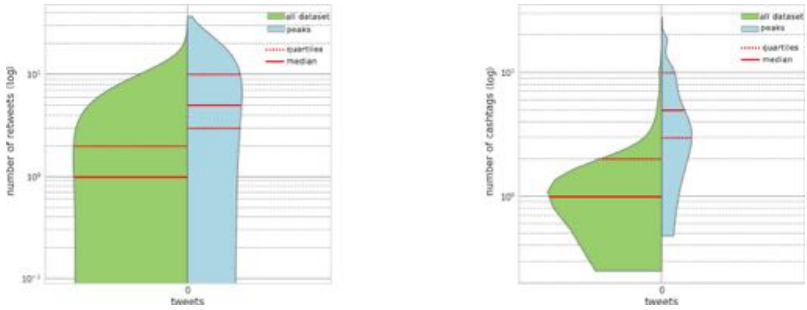

(a) Retweets per tweet. (b) Cashtags per tweet.

Fig. 7. Beanplots showing the differences in the number of retweets per tweet and in the number of cashtags per tweet, for all tweets of the dataset (green-colored) and for peak tweets (light-blue-colored). Peak tweets feature a higher number of retweets and a higher number of cashtags per tweet.

5 ANATOMY OF FINANCIAL SPAM

In this section, we evaluate different hypotheses to thoroughly understand the reasons why so many seemingly unrelated cashtags co-occur in peak tweets and the reason for the high percentage of retweets in peaks.

5.1 Visualizing Co-occurring Stocks

We begin by computing and visualizing the graph of co-occurring stocks for our whole dataset and by comparing it with the graph of stocks that co-occur only in peak tweets. Our co-occurrence graphs represent the collective interconnections of stocks based on their paired presence within tweets. For the sake of clarity, graphs in Figures 8 and 9 only show stocks whose degree $\geq$95.

Figure 8 shows the co-occurrence graph of stocks mentioned in all tweets of our dataset. Stocks are colored according to their market. As shown, the core of the graph is mainly composed of stocks belonging to NASDAQ (blue-colored) and NYSEARCA (green-colored) markets. In addition to stocks of the five markets already introduced, Figure 8 also shows cashtags related to cryptocurrencies (yellow-colored). This is because, in Twitter, cryptocurrencies are labeled with cashtags, similarly to stocks. However, cryptocurrencies are not traded in regulated financial markets and, hence, in Figure 8, they are labeled as OTHERS. As shown, cryptocurrencies represent a large cluster of our graph, with a few highly important nodes such as *Bitcoin* ($BTC) and *Ethereum* ($ETH). Quite intuitively, cryptocurrencies are, however, well separated from the rest of the graph, meaning that they rarely co-occur in tweets with stocks traded in financial markets. Finally, in Figure 8, OTCMKTS stocks (red-colored) cover only a small and peripheral portion of the graph.

In Figure 9, we recreate the co-occurrence graph by only considering peak tweets. This time, the core of the graph is mainly composed of stocks from NASDAQ (blue-colored) and OTCMKTS (red-colored). More precisely, OTCMKTS stocks are not in the periphery of the graph, but instead are well interconnected with NASDAQ stocks. In addition, the degree of many OTCMKTS stocks is comparable to that of NASDAQ stocks. Intuitively, this means that OTCMKTS stocks appear very frequently in peak tweets, and that they often co-occur in such tweets with NASDAQ stocks.

5.2 Analysis of Co-occurring Stocks by Industrial Classification

Previous work has investigated the co-occurrences of stocks in weblogs and their relation to real-world events. In particular, authors of Reference [51] applied a clustering technique over a stock co-occurrences matrix, identifying a number of clusters containing highly correlated stocks. Results

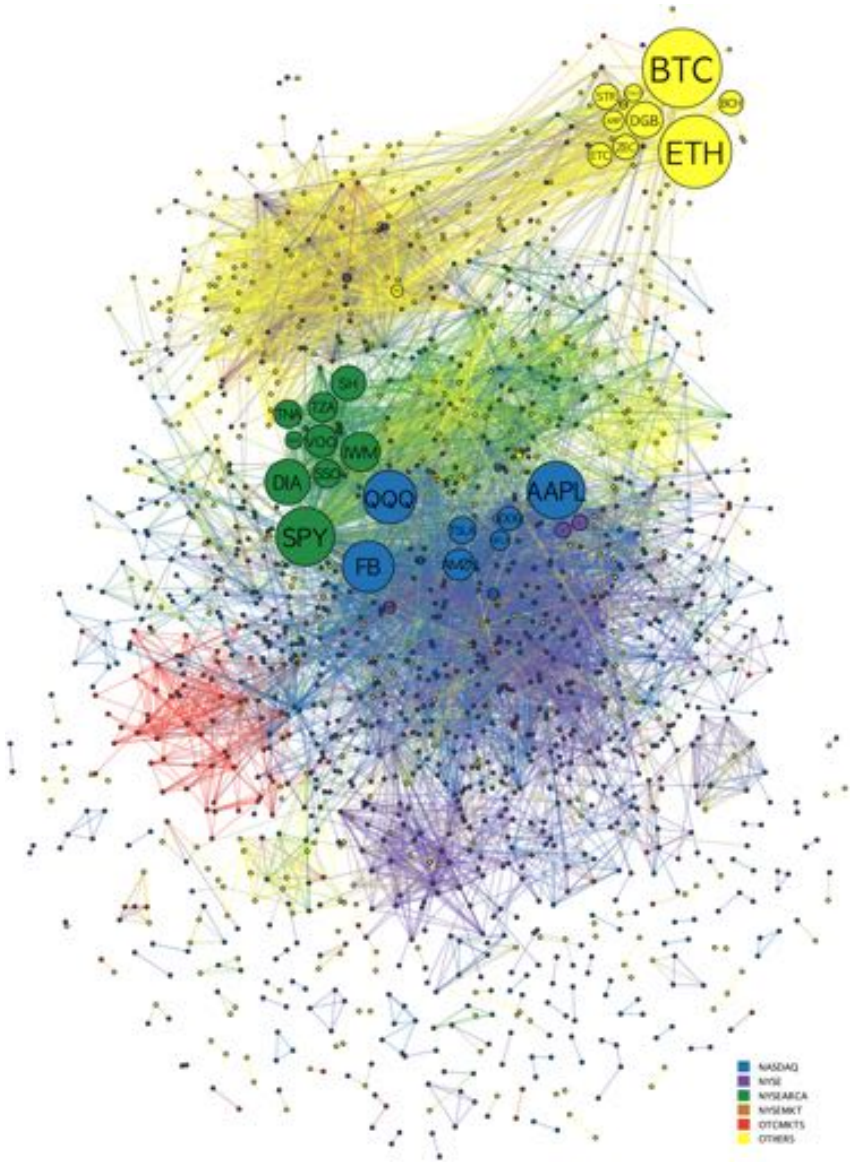

Fig. 8. Co-occurrence graph of stocks mentioned in all tweets of our dataset. OTCMKTS stocks (red-colored) hold a peripheral position in the graph.

of this study highlighted that stocks that co-occur in blog articles as a consequence of real-world events belong to the same industrial sector. In other words, results of Reference [51] support the assumption that stocks that legitimately appear related to one another in weblogs (or microblogs) are also related in real-world. Thus, as a consequence of common sense and previous studies, it would be suspicious for some stocks to appear related (i.e., co-occurring) in microblogs, without being related (i.e., belonging to the same industrial sector) in real-world.

To evaluate whether co-occurring stocks in peak tweets of our dataset are also related in real-world, we exploited the TRBC classification previously introduced in Section 3.2. Specifically, for

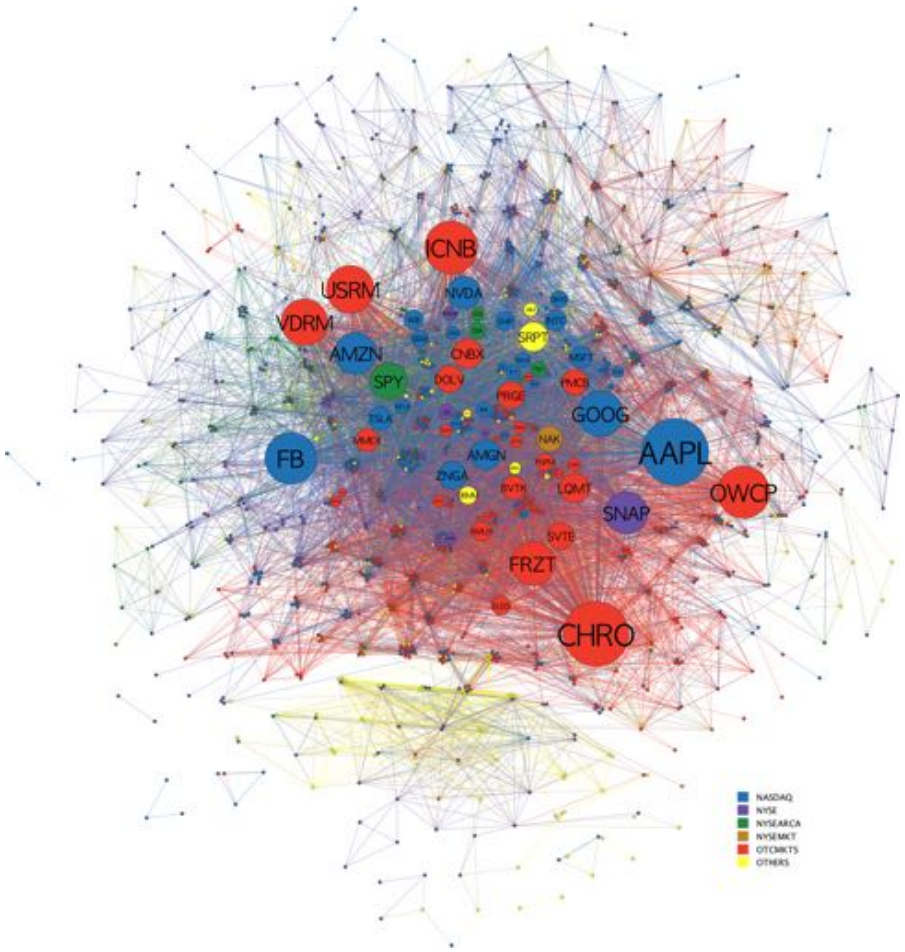

Fig. 9. Co-occurrence graph of stocks mentioned in peak tweets. OTCMKTS stocks (red-colored) are central to the graph and are strongly connected with a number of NASDAQ stocks (blue-colored).

each tweet $t \in \mathbf{t}$, we measured the extent to which the stocks mentioned in t belong to the same (or to different) TRBC class(es), for all the 5 hierarchical levels of TRBC. As a measurement for the difference in TRBC classes across stocks in a tweet, we leveraged the notion of *entropy*. Thus, given a tweet $t \in \mathbf{t}$ containing X distinct cashtags (i.e., each one associated with a different company) and the level j of TRBC with N_j classes, we first built the list of TRBC classes of the X companies mentioned in t:

$$\mathbf{c} = (c_1, c_2, \ldots, c_X).$$

Then, we computed the normalized Shannon entropy of the TRBC classes in $\mathbf{c}$, for TRBC level j, as:

$$H^{\mathbf{c}}_{\text{norm}}(j) = \frac{-\sum_{i=1}^{N_j} p^{\mathbf{c}}_i \log_2 p^{\mathbf{c}}_i}{H_{\max}(j)},$$

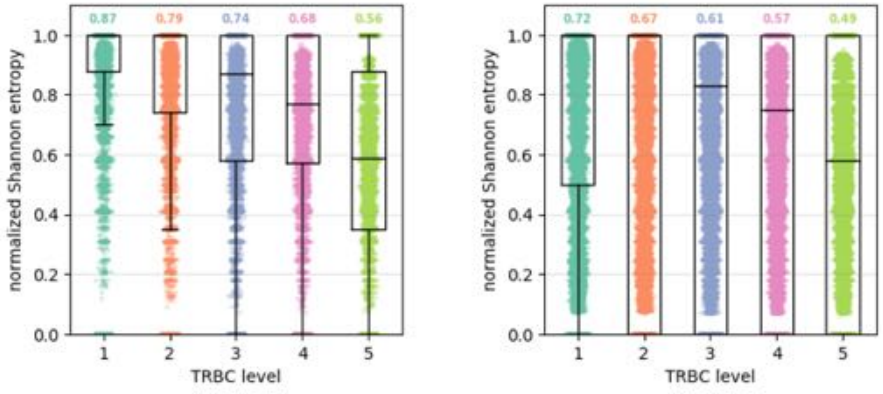

(a) Peaks tweets. (b) All dataset.

Fig. 10. Normalized Shannon entropy of the industrial (TRBC) classes of co-occurring stocks in tweets. TRBC level 1 has the finest grain, while level 5 has the coarsest grain. As shown, median entropy >0.5 for all 5 TRBC levels, meaning that co-occurring companies in tweets are largely unrelated. Mean entropy values are reported above the boxplot and scatterplot distributions. Mean entropy measured for peak tweets is always higher than that measured for all tweets of the dataset. All differences are statistically significant.

where $p_i^{\mathbf{c}}$ is the empirical probability that TRBC class i appears in $\mathbf{c}$, and $H_{\max}(j)$ is the maximum theoretical entropy for TRBC level j:

$$H_{\max}(j) = -\log_2 \frac{1}{X}.$$

Because of the normalization term, $0 \leq H^{\mathbf{c}}_{\text{norm}} \leq 1$. Thus, $H^{\mathbf{c}}_{\text{norm}} \sim 0$ implies companies of the same industrial sector, while $H^{\mathbf{c}}_{\text{norm}} \sim 1$ implies unrelated companies.

Intuitively, considering that the 5 TRBC levels are hierarchical, we expect $H^{\mathbf{c}}_{\text{norm}}$ to be higher (i.e., more heterogeneity) for fine-grained TRBC levels, while we expect $H^{\mathbf{c}}_{\text{norm}}$ to be lower (i.e., less heterogeneity) for the topmost, coarse-grained TRBC level. Results of this experiment, with TRBC level j ranging from the lowest level 1 to the topmost level 5, are shown in Figure 10(a). For every TRBC level, a boxplot and a scatterplot show the distribution of normalized entropy measured for each peak tweet. As expected, $H^{\mathbf{c}}_{\text{norm}}$ actually lowers when considering coarse-grained TRBC levels, as shown by the median value of the boxplot distributions. Nonetheless, median $H^{\mathbf{c}}_{\text{norm}} > 0.5$ for all 5 TRBC levels, meaning that co-occurring companies in peak tweets are largely unrelated. Figure 10(b) shows the result of the same measurement carried out on all tweets of our dataset, rather than only on peak tweets. Interestingly, the entropy measured in all our dataset is smaller than that measured for peak tweets, for all 5 TRBC levels. In turn, this means that co-occurring companies in peak tweets are overall less related than those co-occurring in tweets not belonging to a peak. Differences between the entropies measured for peak tweets and for all the dataset are statistically significant for all 5 TRBC levels, with all p-values <0.01 according to a two-sample Kolmogorov-Smirnov test.

Notably, even for fine-grained TRBC levels, there is a minority of peak tweets for which we measured $H^{\mathbf{c}}_{\text{norm}} = 0$. These tweets might actually contain mentions to companies related also in real-world.

Summarizing, the results of this experiment seem to suggest that, overall, co-occurrences of stocks in peak tweets are not motivated by the fact that stocks belong to the same industrial or economic sectors.

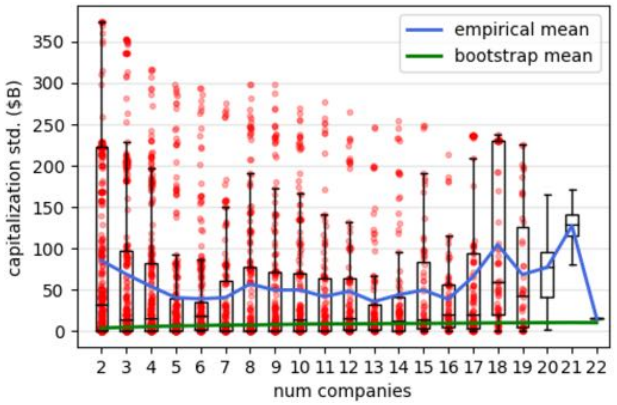

Fig. 11. Standard deviation of the capitalization of co-occurring companies in peak tweets, and comparison with a bootstrap. The large measured standard deviation implies that high-cap companies co-occur with low-cap ones.

5.3 Analysis of Co-occurring Stocks by Market Capitalization

Since real-world relatedness (as expressed by industrial classification) is not a plausible explanation for co-occurring stocks in our dataset, we now turn our attention to market capitalization. We are interested in evaluating whether a relation exists between the capitalization of co-occurring stocks. For instance, legitimate peak tweets could mention multiple stocks with similar capitalization. Conversely, malicious users could try to exploit the popularity of high-cap stocks by mentioning them together with low-cap ones.

One way to evaluate the similarity (or dissimilarity) in market capitalization of co-occurring stocks is by computing statistical measures of spread, *standard deviation* (std.) being a straightforward one. Thus, for each peak tweet $t \in \mathbf{t}$, we computed the std. of the capitalization of all companies mentioned in t. Results are shown in Figure 11, where boxplots and scatterplots are depicted as a function of the number of distinct companies mentioned in tweets. Then, to understand whether the measured spread in capitalization is due to the intrinsic characteristics of our dataset (i.e., the underlying statistical distribution of capitalization) or to other factors, we compared mean values of our empirical measurements with the result of a *bootstrap*. For bootstrapping the std. of tweets that mention x companies, we randomly sampled 10,000 groups of x companies from our dataset. Then, for each of the 10,000 random groups, we computed the std. of the capitalization of the x companies of the group. Finally, we averaged results over the 10,000 groups. This procedure is executed for $x = 2, 3, \ldots, 22$, thus covering the whole extent of Figure 11.

Results in Figure 11 highlight a large empirical std. between the capitalization of co-occurring companies. This means that in our peak tweets, high-cap companies co-occur with low-cap ones. Moreover, the measured std. is larger than that obtained with the bootstrap. In turn, this means that the large difference in capitalization can not be explained by the intrinsic characteristics of our dataset, but, rather, it is the consequence of an external action.

The previous experiment has already led to interesting results. However, it does not allow to draw insights into the possibly different characteristics of stocks traded in different markets. To evaluate the capitalization of co-occurring stocks, for stocks of different markets, we evaluated the *assortativity* of the co-occurrence graph of stocks mentioned in peak tweets. The graph used for this experiment is the one depicted in Figure 9. The assortativity is computed on the capitalization

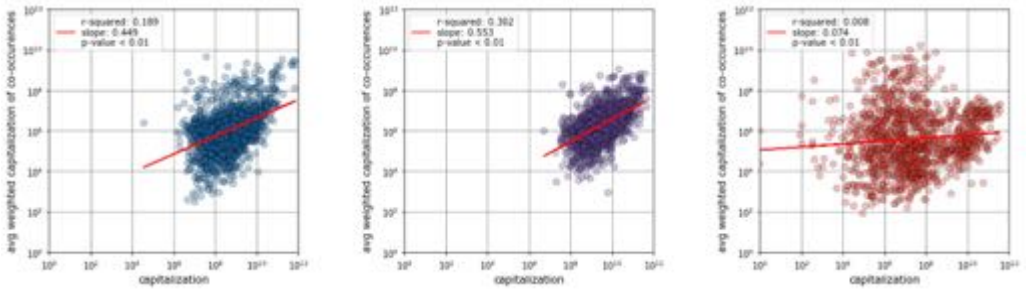

(a) NASDAQ. (b) NYSE. (c) OTCMKTS.

Fig. 12. Assortativity plots for the co-occurrence graph of stocks mentioned in peak tweets. Assortativity is computed out of stock capitalization, and results are grouped by market.

of the nodes (i.e., companies) of the graph, rather than on their degree, as it is typically done with this kind of analysis.

Specifically, for every stock, we compare its capitalization with the weighted mean of the capitalizations of its neighbors in the graph. The weighting factor is based on the number of co-occurrences between stocks (i.e., the weight of the edge in the co-occurrence graph). Results are presented in Figure 12 as scatterplots with a linear fit and are grouped by market. In Figure 12, we only show plots for NASDAQ, NYSE, and OTCMKTS, since they represent the most interesting results. As shown, stocks of NASDAQ and NYSE, the most important markets of our dataset, are assortative (slopes equal 0.44 and 0.55). In other words, high-cap stocks of NASDAQ and NYSE typically co-occur with other high-cap stocks. This behavior is consistent with what one would intuitively expect. Conversely, OTCMKTS stocks are almost all non-assortative, as demonstrated by slope ∼0.

It is also important to note that while the assortativity of NASDAQ and NYSE stocks is higher when considering peak tweets instead of all the tweets of our dataset, for OTCMKTS stocks, we measure the opposite behavior. This means that, in peak tweets, OTCMKTS stocks co-occur with high-cap stocks more often than when considering all our dataset.

5.4 Social and Financial Importance

So far, we demonstrated that tweets responsible for generating peaks mention a large number of unrelated stocks, some of which are high-cap stocks while the others are low-cap ones. Adding to these findings, we are also interested in assessing the relation between the *social* and *financial importance* of our 30,032 stocks. Financial importance of a stock i can be measured by its market capitalization C_i. Social importance can be quantified as the number of times a stock is mentioned in stock microblogs. Intuitively, we expect a positive correlation between stock capitalization and mentions, meaning that high-cap stocks are mentioned more frequently than low-cap stocks. Notably, this positive relation has already been measured in a number of previous works, such as Reference [59], and has been leveraged for predicting stock prices.

By exploiting our data in Table 1, we can make a first assessment of this relation over the whole dataset and compare it with that measured for peak tweets. Specifically, in Table 3, we report the values of two well-known rank correlation measures—namely, Spearman's rank correlation coefficient (ρ) and Kendall's rank correlation coefficient (τ)—between the capitalization of a stock and the number of tweets mentioning that stock. The rank correlation is computed for all stocks of the five markets. When considering all our dataset, for stocks of all markets except OTCMKTS, we find a positive correlation confirming our previous hypothesis. Instead, stocks of OTCMKTS feature negligible rank correlations over all the dataset. Even more interestingly, the significant correlation

Table 3. Rank Correlation Between Market Capitalization and Number of Tweets

	Spearman's ρ		Kendall's τ	
Markets	*all dataset*	*peaks*	*all dataset*	*peaks*
NASDAQ	0.4074	0.0772	0.2960	0.0526
NYSE	0.6347	0.3497	0.4703	0.2452
NYSEARCA	0.4318	0.1429	0.2966	0.1429
NYSEMKT	0.1054	0.0420	0.0719	0.0215
OTCMKTS	0.0778	− 0.2658	0.0556	− 0.1758

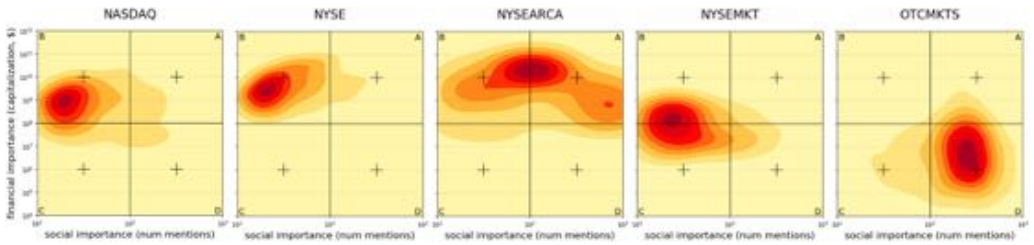

Fig. 13. Kernel density estimation of social and financial importance for stocks of the five considered markets. OTCMKTS stocks have a suspiciously high social importance, despite their low financial importance, in contrast to stocks of all other markets.

measured over all the dataset for stocks of important markets almost completely disappears when only considering peak tweets. Furthermore, OTCMKTS stocks in peak tweets even feature a moderate negative correlation. In other words, results of this experiment imply that the less capitalized stocks in OTCMKTS are more likely to appear in peak tweets than the more capitalized ones, a behavior that is both counterintuitive and contrasts with results of previous works. In turn, this further highlights the presence of suspicious behaviors in peaks.

With the goal of better evaluating the relationship between social and financial importance of stocks appearing in peaks, we also performed an additional experiment as follows: Given a stock i and a peak p, we counted the number of times that i is mentioned in peak tweets of p. We repeated this measurement for every peak p, and we computed the median value of these measurements that represents the social importance of stock i in all peak tweets. Then, for every stock, we plotted its measurement of social importance versus that of financial importance, and we visually grouped stocks by their market. To avoid overplotting, we performed a bivariate (i.e., 2D) kernel density estimation, whose results are shown in Figure 13. For the sake of clarity, we split the social–vs–financial space into four sectors. *Sector A* (top-right) defines a region of space with stocks having both a high social and financial importance. Stocks in *Sector B* (top-left) are characterized by high financial importance but low social importance. Stocks in *Sector C* (bottom-left) have both low social and financial importance, while stocks in *Sector D* (bottom-right) have high social importance despite low financial importance.

By comparing stock densities of different markets in Figure 13, we see that OTCMKTS stocks almost completely lie in *Sector D*. All other markets have their stock densities lying mainly in *Sector B* and *Sector A*. In other words, OTCMKTS stocks have a suspiciously high social importance (i.e., they are mentioned in many tweets and across many peaks), despite their low financial importance. Results for all other markets are more intuitive, with NYSEARCA stocks achieving the best combination of social and financial importance. Summarizing, we measured a positive

relation between social and financial importance when considering all stock microblogs shared during the five months of our study. However, when focusing our analysis on peaks in stock microblogs, we observed a suspicious behavior related to OTCMKTS stocks.

6 ANALYSIS OF SUSPICIOUS USERS

In previous sections, we identified a wide array of suspicious phenomena related to stock microblogs. In particular, peaks in microblog conversations about high-cap stocks are filled with mentions of low-cap (mainly OTCMKTS) stocks. Such mentions can not be explained by real-world stock relatedness. Moreover, the peaks in microblog conversations are largely caused by mass retweets. Despite not having been studied before, this scenario resembles those recently discovered when investigating the activities of bots tampering with social political discussions [28, 37, 63]. Unfortunately, systems for automatically detecting spam in stock microblogs have yet to be developed. However, recent scientific efforts have led to the development of several general-purpose bot and spam detection systems.

6.1 Digital DNA for Social Bot Detection

In this section, we employ a state-of-the-art bot and spam detection system, specifically developed for spotting malicious group activities to classify suspicious users [26, 29]. The goal of this experiment is to assess whether users that shared/retweeted the suspicious peak tweets we previously identified are classified as bots. In turn, this would bring definitive evidence of bot activities in the stock microblogs that we analyzed. The system in References [26, 29] performs bot detection in two steps. First, it encodes the online behavior of a user into a string of characters that represents the *digital DNA* (DDNA) of the user. Then, multiple DDNA sequences—one for each user of the group under investigation—are compared to one another by means of string mining and bioinformatics algorithms. The system classifies as bots those users that have suspiciously high similarities among their DDNA sequences. Notably, the system in References [26, 29] proved capable of accurately detecting also "evolved" bots ($F1 = 0.97$), such as those described in Reference [36].

Because of the computationally intensive analyses performed by References [26, 29], we constrained this experiment to the 100 largest peaks (i.e., those generated by the greatest number of tweets) of our dataset. It is worth noting, however, that, because of the heavy-tailed distribution of tweets per peak, constraining our analysis to the top-100 peaks does not undermine the generalizability of our results. Indeed, the 100 largest peaks contain the most prominent part of the dataset: the 44.30% of tweets in all peaks, posted by the 46.67% of the users who contributed to all the peaks. In other words, constraining the analysis to the 100 largest peaks actually means considering almost 50% of the whole dataset, both in terms of tweets and users.

Starting from those top-100 peaks, we then analyzed the 25,957 distinct users that shared or retweeted at least one peak tweet. Behavioral information needed by the detection system to perform user classification have been collected by crawling the Twitter timelines of such 25,957 users. Notably, the bot-detection system classified as much as 18,509 (71%) of the analyzed users as bots. Our results contrast those reported in Reference [74], in which authors indicate that the bot population ranges between 9% and 15% of all active accounts on Twitter. Our finding suggests that, within the domain of Twitter stock-related messages, the presence of bots among the most active accounts is higher than that measured on average. In turn, this raises concerns about the reliability of stock-related messages in Twitter and highlights the need for targeted bot and financial spam detection techniques. This large presence of bots could be explained by the assumption that the average users are not experts—and possibly not even interested—in finance, and therefore there are fewer people who are actively sharing stock-related messages. Thus, bots represent a larger

Fig. 14. Examples of suspicious users classified as bots. The many characteristics shared between all these users (e.g., name, profile picture, social links) support the hypothesis that they are part of a larger botnet.

Fig. 15. Example tweets from two suspicious users classified as bots.

share of the active Twitter population in the financial domain, with respect to the overall Twitter population.

Figure 14 shows six examples of users classified as bots, while Figure 15 shows some tweets of the same users. To understand the characteristics of these newly uncovered financial bots, we manually inspected a sample of those users classified as bots to determine the nature and the characteristics of their accounts. This process allowed to better understand the results of the DDNA technique and to identify characteristics shared between all the users (e.g., similar name, join date, profile picture, etc.), supporting the hypothesis that they are part of a larger botnet. Moreover, we also performed an automated analysis of all the bots to understand their behavioral characteristics. Interestingly, we found no significant difference between bots and humans in the use of different types of Twitter entities (e.g., hashtags, images, cashtags, etc.). In other words, the content of the messages shared by bots does not appear to be significantly different from the content of human-generated messages. Instead, we noticed that bots have a lower number of followers and friends (followees) with respect to humans. In addition, bots also tend to interact more often than humans, especially via retweets.[11] This outcome supports our previous results related to the high retweeting activity of bots, thus explaining the large number of retweets in our peaks and among `OTCMKTS` stock microblogs. In turn, these findings could be used in the future as features for discriminating between bots and humans.

Retweets have already been associated with frauds in a number of previous works [28, 40, 56], although mass, synchronized retweets had never been reported for stock microblogs. The widespread use of retweets for malicious purposes is due to the fact that retweeting other content is a very simple and elementary activity that does not require the creation of a new message from scratch. In addition, retweets increase the number of users exposed to a specific piece of content,

[11] Complete details of these analyses, including figures, are available upon request to the authors.

thus enabling bots to quickly reach larger audiences. In the financial domain, many trading algorithms monitor social networks to identify "hot" stocks—i.e., companies that are likely to receive a large amount of attention by investors—ahead of time. As a consequence, artificially inflating the number of retweets containing specific cashtags is a simple, yet effective, way of giving the impression of a large market interest for some specific stocks.

Summarizing, the overwhelming ratio of bots that we discovered among large peaks discussing popular stocks raises serious concerns over the reliability of stock microblogs.

6.2 Twitter Bot Detection

In our previous experiment, we relied on a state-of-the-art bot-detection technique to classify our accounts. Here, following a procedure originally used in Reference [28], we also evaluate whether Twitter itself detected and suspended the suspicious accounts that we identified. In fact, accounts that are suspected of performing malicious activities or that violate Twitter's terms of service get suspended by Twitter.

To carry out this experiment, we exploited Twitter's responses to API calls and, in particular, Twitter error codes.[12] Given a query to a specific account, Twitter APIs reply with information regarding the status of the queried account. API queries to a suspended account result in Twitter issuing *error code 63*. Instead, for accounts that are still active, Twitter replies with the full metadata information of the account, without issuing any error.

Results of this experiment show that, out of the 25,957 suspicious accounts, as many as 9,490 (37%) accounts were suspended by Twitter between November 2017 and May 2018. This result is a clear demonstration that many of the accounts responsible for creating the peaks in financial discussions are actually bots. It is not surprising that the *digital DNA*-based technique [26, 29] detected more bots than Twitter (18,509 versus 9,490). Indeed, it has been recently demonstrated that state-of-the-art detection techniques are more effective than Twitter at detecting sophisticated bots [28]. Moreover, to avoid closing accounts of legitimate users by mistake, Twitter is typically conservative with its suspension policy. Finally, there is a very large overlap between the accounts suspended by Twitter and those labeled as bots via the DDNA technique: 8,887 out of 9,490 accounts (~94% of all Twitter suspensions).

6.3 Baselines Comparison

Furthermore, we also compared the results from our DDNA-detection technique with three intuitive baselines. The first baseline consists in selecting a threshold T_1, such that suspicious users with a *retweet rate* (i.e., ratio of retweets over all posted tweets) $>T_1$ are classified as bots. To select this threshold, we used a simple statistical technique called *box plot rule*, already adopted in previous spam and bot detection works [75] and defined as follows: Let $Q_{1,1}$ and $Q_{1,3}$ be the lower and upper quartile, respectively, for retweet rates of a set of genuine (human-operated) accounts. A retweet rate value is an outlier if it lies beyond the *upper inner fence*: $Q_{1,3} + 1.5 \times (Q_{1,3} - Q_{1,1})$. Thus, we selected the upper inner fence of the distribution as the threshold T_1. Since this *anomaly-detection* technique requires a dataset of legitimate accounts, we computed the value for this threshold on the human datasets described in References [25, 28]. Table 4 shows the results of this comparison as a confusion matrix. The baseline agrees with our DDNA bot classifier for 87% of the accounts. However, it appears as less forgiving than our technique, as shown by the 3,305 accounts that we labeled as human but that were classified as bots by the *retweet rate* baseline. In any case, results of this comparison suggest that an exceptionally high number of retweets performed by an account is a strong sign of automated behavior.

[12] https://developer.twitter.com/en/docs/basics/response-codes.html.

Table 4. Baseline 1: Retweet Rate

	baseline 1	
DDNA	*bot*	*human*
bot	18,509	0
human	3,305	4,174

Table 5. Baseline 2: Distinct Cashtags

	baseline 2	
DDNA	*bot*	*human*
bot	6,157	12,352
human	331	7,148

Table 6. Baseline 3: Distinct Cashtag Rates (Ratio of Distinct Cashtags Over All Posted Tweets)

	baseline 3	
DDNA	*bot*	*human*
bot	6,494	12,015
human	15	7,464

The second baseline consists in selecting a threshold T_2, such that suspicious users that used a number of distinct cashtags $>T_2$ are classified as bots. To select a value for this threshold, we simply used the upper quartile of the distribution of distinct cashtags values for all users. Notably, for this baseline, we could not rely on the datasets of References [25, 28], since users from those datasets did not use cashtags and were not involved in stock-related discussions. Thus, they could not serve as a sound ground-truth for comparison. Table 5 shows the comparison for this baseline, which appears as rather ineffective in detecting the bots (agreement with DDNA = 51%). Interestingly, a large number of accounts labeled as bots by digital DNA have been instead labeled as humans by the baseline.

The third baseline consists in selecting a threshold T_3, such that suspicious users with a distinct cashtags rate (i.e., ratio of distinct cashtags over all posted tweets) $>T_3$ are classified as bots. Similarly to the previous baseline, to select a value for this threshold, we simply used the upper quartile of the distribution of distinct cashtags rates for all users. Table 6 reports comparison results for this baseline. Results are similar with respect to the second baseline, in that the agreement with our technique is moderate (54%) and that many accounts labeled as bots by our DDNA technique are instead labeled as human by the baseline.

In conclusion, the retweeting activity of an account is confirmed to be a promising feature for classifying these bots. Conversely, statistics related to the usage of cashtags does not appear to provide useful information.

7 DISCUSSION

Results of our extensive investigation highlighted the presence of spam and bot activity in stock microblogs. For the first time, we described an advertising practice that we called *cashtag piggybacking*, where many financially unimportant (low-cap) stocks are massively mentioned in microblogs together with a few financially important (high-cap) stocks. Analyses of suspicious users suggest that the advertising practice is carried out by large groups of coordinated social bots. Considering the already-demonstrated relation between social and financial importance [59], a possible outcome expected by perpetrators of this advertising practice is the increase in financial importance of the low-cap stocks by exploiting the popularity of high-cap ones.

The potential negative consequences of this new form of financial spam are manifold. On the one hand, unaware investors (e.g., noise traders) could be lured into believing that the social importance of promoted stocks has a basis in reality. On the other hand, the multitude of automatic

trading systems that feed on social information could be tricked into buying low-value stocks. Market collapses such as the *Flash Crash* or disastrous investments such as that of *Cynk Technology* could occur again in the future, with dire consequences. For this reason, a favorable research avenue for the future involves quantifying the impact of social bots and microblog financial spam in stock-price fluctuations, similar to what has already been done at the dawn of financial email spam.

To the best of our knowledge, this is the first exploratory study on the presence of spam and bot activity in stock microblogs. As such, future works related to the characterization and detection of financial spam in microblogs are desirable. Indeed, no automatic system for the detection of financial spam in microblogs has been developed to date. To overcome this limitation, in our analyses we employed a general-purpose bot-detection system. However, such an approach hardly scales on the massive number of users, both legitimate and automated, involved in financial discussions on microblogs. Hence, another promising direction of research involves the development of tools and techniques for promptly detecting promoted stocks, thus avoiding the need for cumbersome user classification. In fact, in light of these findings and starting from the results reported in our study, future works can be focused on the development of new specific detection techniques, tailored to the characteristics of financial spambots to assure that deceptions are detected as early as possible. To this end, some of our most interesting results are related to the high retweeting activity of accounts labeled as bots and their low degree (both inbound and outbound) in the social network. Conversely, characteristics of the content of bot tweets do not seem to provide useful information for an account classification task. For the near future, we thus envision the possibility to exploit these, and other, findings as features to machine-learning financial spam and bot classifiers.

In addition, a strict characterization of the social bots involved in *cashtag piggybacking* spam campaigns (e.g., their behavior and network characteristics) is also needed, to understand whether the synchronized activity of financial botnets are correlated to, or can influence, the stock market's performance or macroeconomic stability.

Finally, we believe it is useful—and worrying at the same time—to demonstrate the presence of bot activity in stock microblogs. Finance, thus, is added to the growing list of domains recently tampered with by social bots—joining the political, social, and commercial domains, to name but a few.

8 CONCLUSIONS

Motivated by the widespread presence of social bots, we carried out the first large-scale, systematic analysis on the presence and impact of spam and bot activity in stock microblogs. By cross-checking 9M stock microblogs from Twitter with financial information from Google Finance, we uncovered a malicious practice aimed at promoting low-value stocks by exploiting the popularity of high-value ones. In these so-called *cashtag piggybacking* spam campaigns, many stocks with low market capitalization, mainly traded in `OTCMKTS`, are mentioned in microblogs together with a few high-capitalization stocks traded in `NASDAQ` and `NYSE`. We showed that such co-occurring stocks are not related by economic or industrial sector. Moreover, the large discussion spikes about low-value stocks are due to mass, synchronized retweets. Finally, an analysis of retweeting users classified 71% of them as bots, and 37% of them were subsequently suspended by Twitter.

Given the severe consequences that this new form of financial spam could have on unaware investors as well as on automatic trading systems, our results call for the prompt adoption of spam and bot detection techniques in all applications and systems that exploit stock microblogs.

SUPPLEMENTARY MATERIALS

Supplementary materials are available in the online version of this article.

REFERENCES

[1] Luca Maria Aiello, Martina Deplano, Rossano Schifanella, and Giancarlo Ruffo. 2012. People are strange when you're a stranger: Impact and influence of bots on social networks. In *Proceedings of the 6th International Conference on Web and Social Media (ICWSM'12)*. AAAI.

[2] Leman Akoglu, Mary McGlohon, and Christos Faloutsos. 2010. Oddball: Spotting anomalies in weighted graphs. In *Proceedings of the 14th Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD'10)*. Springer, 410–421.

[3] Jon-Patrick Allem and Emilio Ferrara. 2018. Could social bots pose a threat to public health? *Amer. J. Pub. Health* 108, 8 (2018), 1005.

[4] Omar Alonso and Kartikay Khandelwal. 2014. Kondenzer: Exploration and visualization of archived social media. In *Proceedings of the 30th International Conference on Data Engineering (ICDE'14)*. IEEE, 1202–1205.

[5] Jalal S. Alowibdi, Ugo A. Buy, S. Yu Philip, and Leon Stenneth. 2014. Detecting deception in online social networks. In *Proceedings of the 6th International Conference on Advances in Social Networks Analysis and Mining (ASONAM'14)*. IEEE/ACM, 383–390.

[6] Marco Avvenuti, Stefano Cresci, Mariantonietta N. La Polla, Carlo Meletti, and Maurizio Tesconi. 2017. Nowcasting of earthquake consequences using big social data. *IEEE Internet Comput.* 21, 6 (2017), 37–45.

[7] Satya Badri, Kyumin Lee, Dongwon Lee, Thanh Tran, and Jason Jiasheng Zhang. 2016. Uncovering fake likers in online social networks. In *Proceedings of the 25th International on Conference on Information and Knowledge Management (CIKM'16)*. ACM, 2365–2370.

[8] Marco T. Bastos and Dan Mercea. 2017. The Brexit botnet and user-generated hyperpartisan news. *Soc. Sci. Comput. Rev.* (2017), 0894439317734157.

[9] Alessandro Bessi and Emilio Ferrara. 2016. Social bots distort the 2016 US presidential election online discussion. *First Mon.* 21, 11 (2016).

[10] Alex Beutel, Wanhong Xu, Venkatesan Guruswami, Christopher Palow, and Christos Faloutsos. 2013. Copycatch: Stopping group attacks by spotting lockstep behavior in social networks. In *Proceedings of the 22nd International Conference on World Wide Web (WWW'13)*. ACM, 119–130.

[11] Johan Bollen, Huina Mao, and Alberto Pepe. 2011. Modeling public mood and emotion: Twitter sentiment and socio-economic phenomena. In *Proceedings of the 5th International Conference on Web and Social Media (ICWSM'11)*. AAAI, 450–453.

[12] Johan Bollen, Huina Mao, and Xiaojun Zeng. 2011. Twitter mood predicts the stock market. *J. Comput. Sci.* 2, 1 (2011), 1–8.

[13] Florian Brachten, Stefan Stieglitz, Lennart Hofeditz, Katharina Kloppenborg, and Annette Reimann. 2017. Strategies and influence of social bots in a 2017 German state election—A case study on Twitter. In *Proceedings of the 28th Australasian Conference on Information Systems (ACIS'17)*.

[14] David A. Broniatowski, Amelia M. Jamison, SiHua Qi, Lulwah AlKulaib, Tao Chen, Adrian Benton, Sandra C. Quinn, and Mark Dredze. 2018. Weaponized health communication: Twitter bots and Russian trolls amplify the vaccine debate. *Amer. J. Pub. Health* 108, 10 (2018), 1378–1384.

[15] Armir Bujari, Marco Furini, and Nicolas Laina. 2017. On using cashtags to predict companies' stock trends. In *Proceedings of the 14th IEEE Annual Consumer Communications & Networking Conference (CCNC'17)*. IEEE, 25–28.

[16] Lorenzo Cazzoli, Rajesh Sharma, Michele Treccani, and Fabrizio Lillo. 2016. A large-scale study to understand the relation between Twitter and financial market. In *Proceedings of the 3rd European Network Intelligence Conference (ENIC'16)*. IEEE, 98–105.

[17] Diego Ceccarelli, Francesco Nidito, and Miles Osborne. 2016. Ranking financial tweets. In *Proceedings of the 39th International Conference on Research and Development in Information Retrieval (SIGIR'16)*. ACM, 527–528.

[18] Nikan Chavoshi, Hossein Hamooni, and Abdullah Mueen. 2016. DeBot: Twitter bot detection via warped correlation. In *Proceedings of the 16th International Conference on Data Mining (ICDM'16)*. IEEE, 817–822.

[19] Nikan Chavoshi, Hossein Hamooni, and Abdullah Mueen. 2016. Identifying correlated bots in Twitter. In *Proceedings of the 8th International Conference on Social Informatics (SocInfo'16)*. Springer, 14–21.

[20] Nikan Chavoshi, Hossein Hamooni, and Abdullah Mueen. 2017. On-demand bot detection and archival system. In *Proceedings of the 26th International Conference on World Wide Web Companion (WWW'17 Companion)*. ACM, 183–187.

[21] Hailiang Chen, Prabuddha De, Yu Jeffrey Hu, and Byoung-Hyoun Hwang. 2014. Wisdom of crowds: The value of stock opinions transmitted through social media. *Rev. Financ. Studies* 27, 5 (2014), 1367–1403.

[22] Eric M. Clark, Chris A. Jones, Jake Ryland Williams, Allison N. Kurti, Mitchell Craig Norotsky, Christopher M. Danforth, and Peter Sheridan Dodds. 2016. Vaporous marketing: Uncovering pervasive electronic cigarette advertisements on Twitter. *PLoS One* 11, 7 (2016), e0157304.

[23] Keith Cortis, André Freitas, Tobias Daudert, Manuela Huerlimann, Manel Zarrouk, Siegfried Handschuh, and Brian Davis. 2017. Semeval-2017 task 5: Fine-grained sentiment analysis on financial microblogs and news. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval'17)*. 519–535.

[24] Stefano Cresci. 2018. *Harnessing the social sensing revolution: Challenges and opportunities*. Ph.D. dissertation. University of Pisa, Pisa, Italy.

[25] Stefano Cresci, Roberto Di Pietro, Marinella Petrocchi, Angelo Spognardi, and Maurizio Tesconi. 2015. Fame for sale: Efficient detection of fake Twitter followers. *Dec. Supp. Systems* 80 (2015), 56–71.

[26] Stefano Cresci, Roberto Di Pietro, Marinella Petrocchi, Angelo Spognardi, and Maurizio Tesconi. 2016. DNA-inspired online behavioral modeling and its application to spambot detection. *IEEE Intell. Systems* 31, 5 (2016), 58–64.

[27] Stefano Cresci, Roberto Di Pietro, Marinella Petrocchi, Angelo Spognardi, and Maurizio Tesconi. 2017. Exploiting digital DNA for the analysis of similarities in Twitter behaviours. In *Proceedings of the 4th IEEE International Conference on Data Science and Advanced Analytics (DSAA'17)*. IEEE, 686–695.

[28] Stefano Cresci, Roberto Di Pietro, Marinella Petrocchi, Angelo Spognardi, and Maurizio Tesconi. 2017. The paradigm-shift of social spambots: Evidence, theories, and tools for the arms race. In *Proceedings of the 26th International Conference on World Wide Web Companion (WWW'17 Companion)*. ACM, 963–972.

[29] Stefano Cresci, Roberto Di Pietro, Marinella Petrocchi, Angelo Spognardi, and Maurizio Tesconi. 2018. Social fingerprinting: Detection of spambot groups through DNA-inspired behavioral modeling. *IEEE Trans. Depend. Secure Comput.* 15, 4 (2018), 561–576.

[30] Stefano Cresci, Fabrizio Lillo, Daniele Regoli, Serena Tardelli, and Maurizio Tesconi. 2018. $FAKE: Evidence of spam and bot activity in stock microblogs on Twitter. In *Proceedings of the 12th International Conference on Web and Social Media (ICWSM'18)*. AAAI, 580–583.

[31] Stefano Cresci, Salvatore Minutoli, Leonardo Nizzoli, Serena Tardelli, and Maurizio Tesconi. 2019. Enriching digital libraries with crowdsensed data. In *Proceedings of the 15th Italian Research Conference on Digital Libraries (IRCDL'19)*. Springer, 144–158.

[32] Stefano Cresci, Marinella Petrocchi, Angelo Spognardi, and Stefano Tognazzi. 2019. On the capability of evolved spambots to evade detection via genetic engineering. *Online Soc. Netw. Media* 9 (2019), 1–16.

[33] Ronen Feldman. 2013. Techniques and applications for sentiment analysis. *Commun. ACM* 56, 4 (2013), 82–89.

[34] Emilio Ferrara. 2015. Manipulation and abuse on social media. *ACM SIGWEB Newslett.* Spring (2015), 4.

[35] Emilio Ferrara. 2017. Disinformation and social bot operations in the run-up to the 2017 French presidential election. *First Mon.* 22, 8 (2017).

[36] Emilio Ferrara, Onur Varol, Clayton Davis, Filippo Menczer, and Alessandro Flammini. 2016. The rise of social bots. *Commun. ACM* 59, 7 (2016), 96–104.

[37] Emilio Ferrara, Onur Varol, Filippo Menczer, and Alessandro Flammini. 2016. Detection of promoted social media campaigns. In *Proceedings of the 10th International Conference on Web and Social Media (ICWSM'16)*. AAAI, 563–566.

[38] Peter Gabrovšek, Darko Aleksovski, Igor Mozetič, and Miha Grčar. 2017. Twitter sentiment around the earnings announcement events. *PloS One* 12, 2 (2017), e0173151.

[39] Saptarshi Ghosh, Bimal Viswanath, Farshad Kooti, Naveen Kumar Sharma, Gautam Korlam, Fabricio Benevenuto, Niloy Ganguly, and Krishna Phani Gummadi. 2012. Understanding and combating link farming in the Twitter social network. In *Proceedings of the 21st International Conference on World Wide Web (WWW'12)*. ACM, 61–70.

[40] Maria Giatsoglou, Despoina Chatzakou, Neil Shah, Christos Faloutsos, and Athena Vakali. 2015. Retweeting activity on Twitter: Signs of deception. In *Proceedings of the 19th Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD'15)*. Springer, 122–134.

[41] Zafar Gilani, Reza Farahbakhsh, and Jon Crowcroft. 2017. Do bots impact Twitter activity? In *Proceedings of the 26th International Conference on World Wide Web Companion (WWW'17 Companion)*. ACM, 781–782.

[42] Eric Gilbert and Karrie Karahalios. 2010. Widespread worry and the stock market. In *Proceedings of the 4th International Conference on Web and Social Media (ICWSM'10)*. AAAI, 59–65.

[43] Martin Hentschel and Omar Alonso. 2014. Follow the money: A study of cashtags on Twitter. *First Mon.* 19, 8 (2014).

[44] Lu Hong and Scott E. Page. 2004. Groups of diverse problem-solvers can outperform groups of high-ability problem-solvers. In *Proc. Natl. Acad. Sci. USA* 101, 46 (2004), 16385–16389.

[45] Tim Hwang, Ian Pearce, and Max Nanis. 2012. Socialbots: Voices from the fronts. *Interactions* 19, 2 (2012), 38–45.

[46] Meng Jiang, Alex Beutel, Peng Cui, Bryan Hooi, Shiqiang Yang, and Christos Faloutsos. 2016. Spotting suspicious behaviors in multimodal data: A general metric and algorithms. *IEEE Trans. Knowl. Data Eng.* 28, 8 (2016), 2187–2200.

[47] Meng Jiang, Peng Cui, Alex Beutel, Christos Faloutsos, and Shiqiang Yang. 2014. Inferring strange behavior from connectivity pattern in social networks. In *Proceedings of the 18th Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD'14)*. Springer, 126–138.

[48] Meng Jiang, Peng Cui, Alex Beutel, Christos Faloutsos, and Shiqiang Yang. 2016. Catching synchronized behaviors in large networks: A graph mining approach. *ACM Trans. Knowl. Discov. Data* 10, 4 (2016), 35.

[49] Meng Jiang, Peng Cui, Alex Beutel, Christos Faloutsos, and Shiqiang Yang. 2016. Inferring lockstep behavior from connectivity pattern in large graphs. *Knowl. Inform. Systems* 48, 2 (2016), 399–428.
[50] Márton Karsai, Kimmo Kaski, Albert-László Barabási, and János Kertész. 2012. Universal features of correlated bursty behaviour. *Sci. Rep.* 2 (2012), 397.
[51] Milad Kharratzadeh and Mark Coates. 2012. Weblog analysis for predicting correlations in stock price evolutions. In *Proceedings of the 6th International Conference on Web and Social Media (ICWSM'12)*. AAAI.
[52] Kai Kupferschmidt. 2017. Bot-hunters eye mischief in German election. *Science* 357, 6356 (2017), 1081–1082.
[53] Sangho Lee and Jong Kim. 2013. Warningbird: A near real-time detection system for suspicious URLs in Twitter stream. *IEEE Trans. Depend. Secure Comput.* 10, 3 (2013), 183–195.
[54] Sangho Lee and Jong Kim. 2014. Early filtering of ephemeral malicious accounts on Twitter. *Comput. Comm.* 54 (2014), 48–57.
[55] Quanzhi Li and Sameena Shah. 2017. Learning stock-market sentiment lexicon and sentiment-oriented word vector from StockTwits. In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL'17)*. 301–310.
[56] Shenghua Liu, Bryan Hooi, and Christos Faloutsos. 2017. HoloScope: Topology-and-spike aware fraud detection. In *Proceedings of the 2017 ACM Conference on Information and Knowledge Management (CIKM'17)*. ACM, 1539–1548.
[57] Xueming Luo and Jie Zhang. 2013. How do consumer buzz and traffic in social media marketing predict the value of the firm? *J. Manag. Inform. Systems* 30, 2 (2013), 213–238.
[58] Xueming Luo, Jie Zhang, and Wenjing Duan. 2013. Social media and firm equity value. *Inform. Systems Res.* 24, 1 (2013), 146–163.
[59] Yuexin Mao, Wei Wei, Bing Wang, and Benyuan Liu. 2012. Correlating S&P 500 stocks with Twitter data. In *Proceedings of the 1st International Workshop on Hot Topics on Interdisciplinary Social Networks Research (SIGKDD'12 Workshops)*. ACM, 69–72.
[60] Symeon Papadopoulos, Kalina Bontcheva, Eva Jaho, Mihai Lupu, and Carlos Castillo. 2016. Overview of the special issue on trust and veracity of information in social media. *ACM Trans. Inform. Systems* 34, 3 (2016), 14.
[61] Neeraj Rajesh and Lisa Gandy. 2016. CashTagNN: Using sentiment of tweets with CashTags to predict stock-market prices. In *Proceedings of the 11th International Conference on Intelligent Systems: Theories and Applications (SITA'16)*. IEEE, 1–4.
[62] Gabriele Ranco, Darko Aleksovski, Guido Caldarelli, Miha Grčar, and Igor Mozetič. 2015. The effects of Twitter sentiment on stock price returns. *PloS One* 10, 9 (2015), e0138441.
[63] Jacob Ratkiewicz, Michael Conover, Mark R. Meiss, Bruno Gonçalves, Alessandro Flammini, and Filippo Menczer. 2011. Detecting and tracking political abuse in social media. In *Proceedings of the 5th International Conference on Web and Social Media (ICWSM'11)*. AAAI, 297–304.
[64] Eduardo J. Ruiz, Vagelis Hristidis, Carlos Castillo, Aristides Gionis, and Alejandro Jaimes. 2012. Correlating financial time series with micro-blogging activity. In *Proceedings of the 5th International Conference on Web Search and Data Mining (WSDM'12)*. ACM, 513–522.
[65] Fabian Schäfer, Stefan Evert, and Philipp Heinrich. 2017. Japan's 2014 general election: Political bots, right-wing internet activism, and prime minister Shinzō Abe's hidden nationalist agenda. *Big Data* 5, 4 (2017), 294–309.
[66] Harald Schoen, Daniel Gayo-Avello, Panagiotis Takis Metaxas, Eni Mustafaraj, Markus Strohmaier, and Peter Gloor. 2013. The power of prediction with social media. *Internet Res.* 23, 5 (2013), 528–543.
[67] Chengcheng Shao, Giovanni Luca Ciampaglia, Alessandro Flammini, and Filippo Menczer. 2016. Hoaxy: A platform for tracking online misinformation. In *Proceedings of the 25th International Conference on World Wide Web Companion (WWW'16 Companion)*. ACM, 745–750.
[68] Jasmina Smailović, Miha Grčar, Nada Lavrač, and Martin Žnidaršič. 2014. Stream-based active learning for sentiment analysis in the financial domain. *Inform. Sci.* 285 (2014), 181–203.
[69] Timm Oliver Sprenger. 2011. TweetTrader.net: Leveraging crowd wisdom in a stock microblogging forum. In *Proceedings of the 5th International Conference on Web and Social Media (ICWSM'11)*. AAAI.
[70] Andrew Tanenbaum and David Wetherall. 2014. *Computer Networks*. 5th Edition. Pearson Education Limited.
[71] Shiliang Tang, Qingyun Liu, Megan McQueen, Scott Counts, Apurv Jain, Heather Zheng, and Ben Y. Zhao. 2017. Echo chambers in investment discussion boards. In *Proceedings of the 11th International Conference on Web and Social Media (ICWSM'17)*. AAAI.
[72] Stefano Tognazzi, Stefano Cresci, Marinella Petrocchi, and Angelo Spognardi. 2018. From reaction to proaction: Unexplored ways to the detection of evolving spambots. In *Proceedings of the 27th Web Conference Companion (WWW'18 Companion)*. ACM, 1469–1470.
[73] Jan van der Tempel, Aliya Noormohamed, Robert Schwartz, Cameron Norman, Muhannad Malas, and Laurie Zawertailo. 2016. Vape, quit, tweet? Electronic cigarettes and smoking cessation on Twitter. *Int. J. Pub. Health* 61, 2 (2016), 249–256.

[74] Onur Varol, Emilio Ferrara, Clayton A. Davis, Filippo Menczer, and Alessandro Flammini. 2017. Online human-bot interactions: Detection, estimation, and characterization. In *Proceedings of the 11th International Conference on Web and Social Media (ICWSM'17)*. AAAI.

[75] Bimal Viswanath, Muhammad Ahmad Bashir, Muhammad Bilal Zafar, Simon Bouget, Saikat Guha, Krishna P. Gummadi, Aniket Kate, and Alan Mislove. 2015. Strength in numbers: Robust tamper detection in crowd computations. In *Proceedings of the 2015 ACM Conference on Online Social Networks (COSN'15)*. ACM, 113–124.

[76] Tianyi Wang, Gang Wang, Bolun Wang, Divya Sambasivan, Zengbin Zhang, Xing Li, Haitao Zheng, and Ben Y. Zhao. 2017. Value and misinformation in collaborative investing platforms. *ACM Trans. Web* 11, 2 (2017), 8.

[77] Fangzhao Wu, Jinyun Shu, Yongfeng Huang, and Zhigang Yuan. 2015. Social spammer and spam message co-detection in microblogging with social context regularization. In *Proceedings of the 24th International on Conference on Information and Knowledge Management (CIKM'15)*. ACM, 1601–1610.

[78] Chao Yang, Robert Harkreader, and Guofei Gu. 2013. Empirical evaluation and new design for fighting evolving Twitter spammers. *IEEE Trans. Inform. Forens. Sec.* 8, 8 (2013), 1280–1293.

[79] Haifeng Yu, Phillip B. Gibbons, Michael Kaminsky, and Feng Xiao. 2010. SybilLimit: A near-optimal social network defense against sybil attacks. *IEEE/ACM Trans. Netw.* 18, 3 (2010), 885–898.

[80] Rose Yu, Xinran He, and Yan Liu. 2015. GLAD: Group anomaly detection in social media analysis. *ACM Trans. Knowl. Discov. Data* 10, 2 (2015), 18.

[81] Yang Yu, Wenjing Duan, and Qing Cao. 2013. The impact of social and conventional media on firm equity value: A sentiment analysis approach. *Dec. Supp. Systems* 55, 4 (2013), 919–926.

[82] Xianchao Zhang, Zhaoxing Li, Shaoping Zhu, and Wenxin Liang. 2016. Detecting spam and promoting campaigns in Twitter. *ACM Trans. Web* 10, 1 (2016), 4.

[83] Ilya Zheludev, Robert Smith, and Tomaso Aste. 2014. When can social media lead financial markets? *Sci. Rep.* 4 (2014), 4213.

Received July 2018; revised December 2018; accepted February 2019