Improving information spread by spreading groups

Alon Sela
Department of Industrial Engineering, Ariel University, Ariel, Israel; Department of Industrial Engineering, Tel Aviv University, Tel Aviv, Israel and Ariel Cyber Innovation Center (ACIC), Ariel University, Ariel, Israel
Orit Milo
Department of Finance, Hebrew University of Jerusalem, Jerusalem, Israel
Eugene Kagan
Department of Industrial Engineering, Ariel University, Ariel, Israel, and
Irad Ben-Gal
Department of Industrial Engineering, Tel Aviv University, Tel Aviv, Israel

Received 20 August 2018
Revised 1 March 2019
14 June 2019
Accepted 12 August 2019

Abstract

Purpose – The purpose of this paper is to propose a novel method to enhance the spread of messages in social networks by "Spreading Groups." These sub-structures of highly connected accounts intentionally echo messages between the members of the subgroup at the early stages of a spread. This echoing further boosts the spread to regions substantially larger than the initial region. These spreading accounts can be actual humans or social bots.
Design/methodology/approach – The paper reveals an interesting anomaly in information cascades in Twitter and proposes the spreading group model that explains this anomaly. The model was tested using an agent-based simulation, real Twitter data and questionnaires.
Findings – The messages of few anonymous Twitter accounts spread on average more than well-known global financial media groups, such as The Wall Street Journal or Bloomberg. The spreading groups (also sometimes called BotNets) model provides an effective mechanism that can explain these findings.
Research limitations/implications – Spreading groups are only one possible mechanism that can explain the effectiveness of spread of tweets from lesser known accounts. The implication of this work is in showing how spreading groups can be used as a mechanism to spread messages in social networks. The construction of spreading groups is rather technical and does not require using opinion leaders. Similar to the case of "Fake News," we expect the topic of spreading groups and their aim to manipulate information to receive growing attention in public discussion.
Practical implications – While harnessing opinion leaders to spread messages is costly, constructing spreading groups is more technical and replicable. Spreading groups are an efficient method to amplify the spread of message in social networks.
Social implications – With the blossoming of fake news, one might tend to assess the reliability of news by the number of users involved in its spread. This heuristic might be easily fooled by spreading groups. Furthermore, spreading groups consisting of a blend of human and computerized bots might be hard to detect. They can be used to manipulate financial markets or political campaigns.
Originality/value – The paper demonstrates an anomaly in Twitter that was not studied before. It proposes a novel approach to spreading messages in social networks. The methods presented in the paper are valuable for anyone interested in spreading messages or an agenda such as political actors or other agenda enthusiasts. While social bots have been widely studied, their synchronization to increase the spread is novel.

Keywords Information spread, Viral marketing, Social networks, Spreading groups, Fake news, BotNets
Paper type Research paper

1. Introduction

Modern social network platforms provide a simple and efficient way to spread messages and increase their influence (Teng *et al.*, 2014) among individuals or communities of users. The results of such message spread (the terms "messages" and "information" are used

Ariel Cyber Innovation Center (ACIC), Ariel University, Ariel, 40700, Israel.

Online Information Review

1468-4527
DOI 10.1108/OIR-08-2018-0245

interchangeably) can vary in scale, from serving the needs of an individual user to having a widespread influence on political (Harvey, 2013), social and economic phenomena. Examples of large-scale social influences where social networks played a major role include the "Arab Spring" (Howard *et al.*, 2011), the "Occupy Wall Street" movement (Bennett and Segelberg, 2011) and the elections in the USA (Allcott and Gentzkow, 2017; Kim and Hastak, 2018). In fact, in several areas, the role of social networks and their smart use by a small group of users can lead to an outcome that can be as influential as that of traditional media.

In the financial markets, the role of social networks as an information spreader is critical. Research has established that the news spread through a social network is often published before it reaches the traditional media (Leskovec *et al.*, 2009). Trust is also a critical aspect, affecting the likelihood of a message's spread (Gorman, 2014). Nevertheless, the blossoming of fake news (Allcott and Gentzkow, 2017; Newman *et al.*, 2017) might suggest that sometimes trust can be misleading, and fake rumors can seem to be trustworthy sources.

Our investigation was motivated by the observation that there are individual Twitter accounts with especially high levels of average retweet rates that are substantially higher than those of well-known global media companies such as The New York Times or The Wall Street Journal. We define "well-known users" to be international media outlets who are traditionally "information broadcasters." More specifically, we define well-known users as international media companies with more than 1m followers. We expect such organizations to play a more significant role in information dissemination compared to a private account or small companies.

To explain these anomalies, we constructed a novel theoretical model that demonstrates the boosting of a message's spread by spreading groups. Such spreading groups are organized sub-structures of interconnected users that intentionally amplify the exposure of a message that they wish to spread. To accomplish this goal, they utilize a ping-pong-like interchange of the message within the group's members in its early spreading stages. Through this strategy of transmitting a message over and over within the group in the early stages of its spread, these groups ignite a larger information cascade, and help the message's spread to other parts of the social network that are external to the group itself at later stages.

The possible benefits gained from using spreading groups are clear. In contrast to other message-spreading methods, such as using opinion leaders, the construction of spreading groups is mainly a technical issue. The formation of spreading groups can be accomplished with real Twitter accounts or with the accounts of bots (computer programs that imitate a social network user). Thus, although the initial construction of spreading groups requires some initial technical skill, its later use to gain influence in a social network is easy.

The paper is organized as follows. Section 2 provides a short history of rumor spread in stock-related domains and suggests a motivation for the creation of spreading groups. Section 3 presents the theoretical model and, using simulations, demonstrates how by using only spreading groups a message can be spread to a larger audience. This section defines the model's characteristics and formulates different criteria for the intentional spread of information. Section 4 describes the methods of Twitter data analysis, which is followed by the experimental results in Section 5. The Results section begins with the theoretical model and then details the Twitter data set analysis and its results. Section 6 discusses the findings, and Section 7 concludes by describing the implications of bot technology in message spread and its relevance to spreading groups.

2. Background

Studies of stock-related news rumors and their influence on financial markets can be traced back to the rise of internet message boards and forums. In early 1998, for example, the correlation between the number of published messages about companies on the Yahoo Finance Board and their shares' values was noticed and reported (US Securities and Exchange Commission, 1998). In their study of the social influence on the stock market,

DeMarzo *et al.* (2003) created a model of network activity and traced its influence. The authors argued that users form their opinions with respect to the opinions of other users. They are swayed more by the appearance of repeated messages than by the accuracy of the information presented in these messages. Similarly, Cao *et al.* (2002) claimed that financial traders might be less affected by the actual values of companies than by extensive conversations about the stock. Based on these findings, Antweiler and Frank (2004) demonstrated that a higher number of message postings predicts negative stock returns, greater volume, and more volatility. This opinion coincides with previous findings (Harris and Raviv, 1993) claiming that opposite opinions among posted messages tend to be followed by increased trading volume. Thus, an increase in the spread of a message, regardless of its content, might have a real financial impact.

With the development of communication technology, message boards were mostly replaced by instant messaging platforms and social networks such as Twitter and WhatsApp. In addition to the ability to accelerate and facilitate communication between users, as well as the growing amounts of irrelevant information that the users receive (Hirshleifer and Teoh, 2003), these platforms enabled the creation of messaging groups that facilitated information spread through customized, digitalized sub-networks of a defined type of knowledge.

However, these studies focused on the flow of rumors in financial markets, investigating the influence of rumors on stock prices, and were not concerned with social networks.

The proposed model of information flow follows the general approach of information cascades that Kempe *et al.* (2003) studied. They analyzed the problem of selecting a correct set of nodes in a graph such that seeding these nodes (i.e. setting these nodes as infected) would maximize the number of infected nodes at the end of a "viral" spreading process. They also defined two generic models that captured much of the essence of information spread through social networks. Their first model is the linear threshold model, and the second is the independent cascade model. Both models assume an initial set of k seeds that are selected prior to the spread. Then, a spreading process is performed until no further activations are possible. The measure assessed for each combination of initial seeds is the final number of infected nodes. The difference between the two models is an indicator of the dynamics of the spreading process. In the linear threshold model, a spread occurs only if the sum of the influences (i.e. the weights) of the infected neighbors is greater than a threshold θ, meaning when $\sum_w b_{v,w} > \theta_v$, with $b_{v,w}$ defined as the influence weights. On the other hand, in the independent cascade model an infected node can infect its neighbors only for a single period after it first becomes infected. The linear threshold model thus captures the essence of social influence, where the more the adopters of a position or idea within one's circle of friends, the more likely one is to adopt that same position or idea. In contrast, the independent cascade model captures the retention loss of human memory, where old news is less likely to spread than novel news.

We implemented a variant of these two models on power law graphs. Power law graphs are commonly used as a first-order approximation for many natural phenomena, including social networks. These graphs describe processes where new user accounts (nodes) are more likely to connect to highly connected nodes than less connected ones (Barabasi and Albert, 1999; Bollobas, 2001). This dynamic process describes the human tendency to connect to more popular, well-connected individuals.

While the network power law structure is an important factor in the spread of information, the micro properties that are required for a single transmission of information between two nodes also need to be addressed. Information spread requires the ability to transmit, as well as the ability to receive and process a new message. The amount of information has grown exponentially during the last 20 years, and people are exposed to amounts of information that have never been seen before. Technologies such as search engines help sort and filter this vast amount of information, but these technologies also have their own biases (Sela *et al.*, 2016).

Taking into consideration that people's attention resources are limited, but the processing requirements are growing exponentially, smaller fractions of information are processed. The limited attention model (Weng *et al.*, 2012) is based on observations related to people's limited capacity for attention, which is usually insufficient to process the entire stream of news that flows in a social network. New information that arrives while one's cognitive resources are occupied (or when people are simply too busy to open their Facebook or Twitter account) are never read nor processed. Given that the speed and quantities of information flow are increasing, but people's cognitive attention capacities remain the same, they are likely to process a decreasing fraction of the messages.

To deal with this issue, many people tend to follow opinion leaders. The study of influencers (opinion leaders) has proposed various centrality measures, which were widely assessed in the early days of network science studies. Numerous centrality measures were proposed, each quantifying a different aspect of influence. Some of the common centrality measures are PageRank, the Eigenvector centrality, the Betweenness centrality, the Kats centrality, and clustering coefficient measures (Newman, 2005). Each of these measures has its own advantages and represents a different facet of influence that characterizes a defined node. Several studies have summarized the vast literature on centrality measures (Kempe *et al.*, 2003; Borgatti, 2005; Aral *et al.*, 2013; Shakarian *et al.*, 2013).

Centrality measures are often used to evaluate and improve seeding policies. Seeds might be selected separately or together as an entire group. Some studies investigated which sets of nodes should be seeded simultaneously to increase the information spread. Other works (Chierichetti *et al.*, 2014a; Michalski *et al.*, 2014) have addressed the question of timing, focusing on determining not just which nodes to seed. Since retention loss reduces the probability of infection as time passes, and since users who receive messages repeatedly tend to get annoyed, there are cases where addressing a central node too early might only decrease the final spreading rates (Sela *et al.*, 2016). Thus, the question of influence maximization should address not only what nodes to seed, but also which nodes to seed and when (Sela *et al.*, 2017; Goldenberg *et al.*, 2018; Sela *et al.*, 2018).

The relationship between the number of exposures to a message and likelihood of retweets has been studied by Zhou *et al.* (2015). They found that the more a person is exposed to a message, the higher the retweet likelihood. These results imply that a connected sub-component in a social network, which is also connected to many naive users, will increase the likelihood that the naive users will retweet a message, since they will be exposed to it more frequently. Thus, there is a clear advantage in creating sub-structures of connected nodes in a social network that are constructed to spread a defined agenda.

These studies establish the ground for the methodology of complex networks and information spread. However, they did not focus on social bots; rather, they were concerned with theoretical modeling and the universal aspects of information spread – regardless if the spreader is human or a bot.

This work adds a novel layer to the study of influence on social networks. It offers an applicable technique, by which the spread of a message can be actively increased by constructing sub-structures of human–bots and using these sub-structures to repeatedly echo a message. While numerous studies (Davis *et al.*, 2016; Echeverria and Zhou, 2017; Ferrara, 2017; Freitas *et al.*, 2015) have investigated the existence of bots in social networks, the synchronized activity of bots by spreading groups has been neglected. Our main contribution is to demonstrate that the spreading group strategy is highly efficient, and that it is probably already used by groups of stock traders.

3. The model

The model compares the infection rates achieved by using a spreading group rather than other approaches to seeding (i.e. selecting initial accounts in a social network and spreading

a message to these accounts). Its results provide insights into the observations obtained from real Twitter messages about the NASDAQ-100 stock market and demonstrate that some rather unknown users reach retweeting rates 10 times higher than those of well-established financial media groups.

Let $D = (V, A)$ be a directed graph (DAG) with n nodes and e edges that represents the Twitter (or similar) social network. The set of vertices $V = \{v_1, v_2, \ldots, v_n\}$ is associated with the individual users in the network, and the set of edges $A = \{e_{ij} = (v_i, v_j) | v_i, v_j \in V\}$ defines the connectivity between the users, where user j follows user i (thus, information spreads from user i to j). The network topology is constructed by the preferential attachment process (Barabasi and Albert, 1999) such that nodes are added to the network one after the other, and the probability that a new vertex $v_j \in V$ connects to a specific existing vertex v_i depends on the relative connectivity between the specific degree k_i and the sum of all existing degrees $\sum_{j=1}^{n} k_j$ as defined in the below equation:

$$p(v_i) = \frac{k_i}{\sum_{j=1}^{n} k_j}, \quad i = 1, 2, \ldots, n. \tag{1}$$

A spreading group of users is denoted by a sub-graph $D' \subseteq D$ that is constructed as follows. Let $D' = (V', A')$, $V' \subset V$ be a subset of vertices of size $\#V' = \lceil nr \rceil$, where $r \in (0, 1)$ is a given ratio value and n is the number of nodes in the entire network. During the preferential attachment construction of D, edges between pairs of vertices in V' are added arbitrarily in such a manner that the average connectivity of the sub-graph D' increases independently of the connectivity of D. More precisely, first (n/2) nodes are created and attached by a preferential attachment BA constructor (Barabasi and Albert, 1999). Then, nr nodes are randomly selected. Third, for each node in V', r additional links are added to connect any of the nodes in V' without self-loops, where we usually use $r = 0.02$. Last, the additional (n/2) nodes are added by a preferential attachment constructor, to grow the network to its full size of n nodes. As a result, the graph D includes a sub-graph D' with greater connectivity. Such a structure corresponds to a network with a distinguishable subset (group) of highly interconnected users.

3.1 Three spreading strategies

As indicated above, the construction of the graph D follows the preferential attachment model and creates a distinguishable sub-graph D' within this graph from which the spreading group might be selected. Now let us consider the initial message seeding and information flow in the network. We denote by $V_s \subset V$ a subset of vertices that are associated with the users who initiate the spread, meaning the message seeds. We consider the following three types of settings for the selection of the vertices in V_s with regard to the selection of nodes in the spreading groups:

(1) The seeded vertices $v \in V_s$ are selected randomly from the set V independently of whether they are in the spreading group V' or not.

(2) The seeding vertices $v \in V_s$ are selected randomly from the set V' of vertices, which is associated with the nodes in the sub-graph D', the spreading group.

(3) The set V_s of seeding vertices is specified as a set of vertices with the highest eigenvector centrality measures (Newman, 2010). In other words, V_s is associated with the group of users who have the strongest influence in the network.

These three methods of selection are illustrated in Figure 1. The selected seeds are marked by red nodes, and the spreading groups' connectivity is marked by an orange sub-graph. The left-hand image plots the first setting for the selection, where seeds are selected

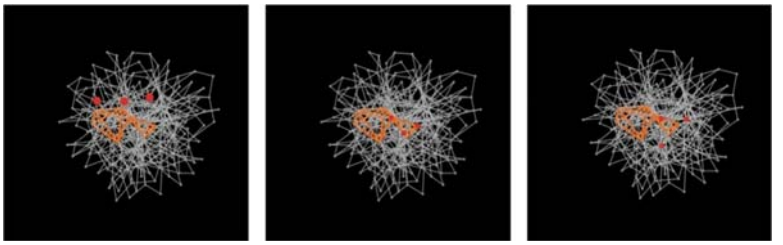

Notes: The spreading group is marked by an orange sub-graph. The three methods of selection are random seed selection (left), random selection from the spreading group (center) and selection of most influential nodes (right)

Figure 1.
Illustration of seed (red nodes) allocation settings

randomly from V. The middle image plots the setting where the seeds are selected randomly from the spreading group V'. In the right-hand image, the seeds are selected by choosing the nodes with the strongest influence, that is, the highest eigenvector centrality, which are not necessarily part of V_s.

After the initial seeding by vertices $v \in V_s$, information is spread according to the independent cascade model (Kempe *et al.*, 2003) with some minor modifications. As in Kempe *et al.*'s (2003) study, the spreading process spreads a message only from a user to his/her neighbors. The receiving user can then spread the message to his/her neighbors with a transmission probability $p(t)$. Unlike the independent cascade model, where a node can transfer a message only in the consecutive period, we used the theory of limited attention (Weng *et al.*, 2012). Therefore, we assumed that the transmission probability decreases over time, but is still possible after the first period that follows the node's infection time. The initial transmission probability for each node during the first period after receiving the message is defined as $p_0 \equiv p(t = 0)$, and is recursively set in the subsequent time to:

$$p(t+1) = \frac{p(t)}{c}, \quad t = 0, 1, 2, 3\ldots \tag{2}$$

In other words, when a user receives a message, the probability of spreading it in the first step ($t = 0$) is set to a maximal value p_0, and then decreases over time. In each time step, a random process defines whether any node v will be infected by its neighbor u, where the probability of such an infection is $p(t)$. When an infection of u succeeds, its time counter is set to $t = 0$. The process of infection can theoretically continue forever, but we define its ending if no users in the graph were infected in two consecutive time steps.

While the dynamics of information flow is rather simple, the comparison of different seeding strategies, particularly the comparison of the seeding results when the spreading group sub-graph is constructed, is the main contribution of the model above.

4. Method

The methods of study included both a simulation and an analysis of a data set collected from Twitter. To analyze the results of the simulation obtained using our model, we utilized NetLogo, an agent-based simulation platform (Wilensky, 1999; Wilensky and Rand, 2015). The simulations used a Barabasi–Albert graph of $n = 10{,}000$ vertices. In these studies, we used the three different seeding methods described in Section 3: seeded vertices $v \in V_s$ with nodes selected randomly from V, seeded vertices $v \in V_s$ with nodes selected randomly from the spreading group and seed nodes starting with the node with the highest eigenvalue centrality measure (influence), then the 2nd, 3ed, …, kth high eigenvalue centrality. The model's parameter space was modified such that the dense cluster V' contained 1, 3 or 5 per cent of the nodes from V.

Second, the number of initial seeds $|V_s|$ was set to 15, 25, or 35 nodes, respectively. Third, the transmission probability p_0 was set to 1%, 5%, or 10 percent. Last, the retention loss factor c was set to 1.5, 3 or 4.5. For each unique combination of parameters, the simulations were run for 20 runs, while recording the final number of infected nodes in the graph in each run.

In parallel to the simulation experiments, we analyzed and compared a data set from Twitter. The raw data included 1,481,444 NASDAQ-100 related tweets that circulated on Twitter throughout June 2014. A filtering process was applied to select tweets that included only the "$" sign before the identifying string. Based on Twitter jargon, such a selection distinguishes between the stock itself (e.g. the $AAPL stock) and other possible objects with the same name (e.g. the fruit or the company that has the name "Apple"). The Twitter jargon contains a few common symbols such as the hashtag symbol #, which represents a term, and the symbol @, which represents a Twitter user's name. The $ symbol is a unique sign in Twitter that is used for stock-related tweets. Thus, filtering tweets that included the $ sign followed by one of the names of any of the 100 stocks in the NASDAQ-100 created a filter for messages related to stock markets only.

To preserve consistency, we will refer to individual tweets sent by a Twitter user as "message" or "tweet," and we will refer to a set of recurrent messages with a similar text as "cascade."

The raw data set included 1.48m tweets, all having a NASDAQ-100 hashtag, that were downloaded by the streaming Twitter API. We first removed all the tweets that were not complete, i.e. not written in English, or had missing/non-readable text. Following this filtering, the data set included 1,131,625 tweets. Then, we grouped these tweets by text, and only kept tweets that appeared more than once. This left the data set with only 27,071 cascades that had at least one repetition. Together with the tweet text, the occurrence record included a field tag indicating the time it was published, the name of the publisher, the location where the tweet was retweeted, and, in cases of retweets, the identifier of the user who initialized the tweet cascade. Since the streaming Twitter API only permits downloading about 1 percent of the entire Twitter data, we could not count on understanding the data structure. In addition, Twitter does not fully reveal the structure of its network. For example, consider the case where person A follows person B, and person B follows person C. Also suppose that a message published by C is retweeted by B, and then B is retweeted again by A. This would appear when downloading the data as a message published by C is retweeted by A, even though C and A do not have a connection, while the connection of B and A will not be revealed. This is the reason why one can sometimes see on Twitter messages by people that one does not follow. If a person that one does follow retweets a message, this message will appear in one's timeline with the name of the original publisher[1].

5. Results

5.1 Simulation results

Given the above-mentioned set of parameters, the simulation study included 4,860 runs (20 runs with each distinct combination of parameters). In each run, the total number of vertices that were eventually exposed to the message by the end of the trial was recorded. The average number of infected vertices at the end of each run, over the trials for different sets of seeding methods, and different number of initial seeds are shown in Figure 2.

Note that if the seeding vertices are selected randomly from the set of vertices, the number of exposed vertices at the end of the simulation is nearly three times smaller than this number when the seeding vertices are selected from the spreading group. Such a result accords with the findings from the evaluation of the Twitter messages presented in the next section, which provide evidence of the existence of such spreading groups in reality. Since the connectivity in the spreading group $G^{'}$ is higher than the connectivity in the network G, the messages in G' circulate more widely and, consequently, are better exposed to the users in

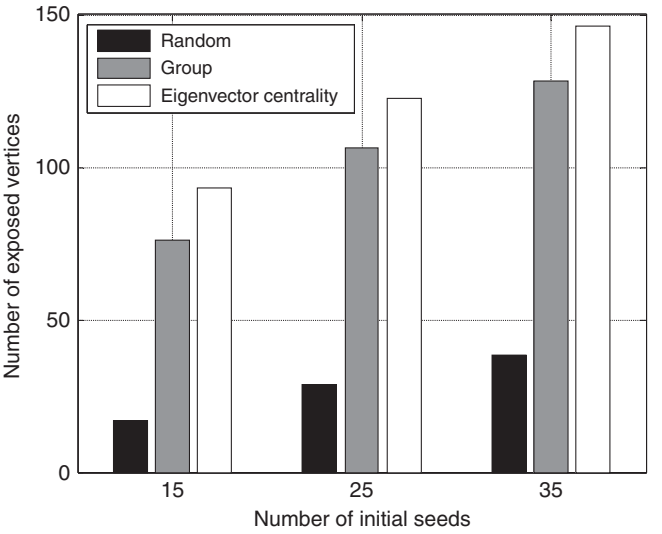

Figure 2.
Average number of exposed vertices for three types of seeding methods and for number of initial seeds

Notes: Black bars correspond to the seeding vertices selected randomly from the set of all vertices, gray bars correspond to the seeding vertices selected from the spreading group and white bars correspond to the seeding of vertices specified by the vertices with the highest Eigenvector centrality measure

the remaining network. In addition, note that the information flow model we used does not consider the meaning or the content of the messages (Harris and Raviv, 1993) or the increased tendency to read a message from different users in the social circle (herd behavior). These influences may highlight the difference between the structures of the random seeds compared to the initiation of the message by the two other seeding groups as described above.

To better understand the mechanism behind the spreading groups, one needs to consider the dynamics of the epidemic spreading model. In epidemic models, such as the SIR model (Funk *et al.*, 2009; Wu *et al.*, 2004), the epidemic is very sensitive in its initial stage of spread. Small efforts at immunization might quash the epidemic. In contrast, in later stages, when the epidemic is already spread into the population, a substantially greater effort would be required to suppress its spread. A similar dynamic exists in the case of information spread. In its echoing of a message in its initial stages of spread, the spreading groups act as a shield that protects and strengthens the spread in its early stages, thus enabling its spread to larger regions of the network at later stages.

Figure 3 illustrates the factors that strengthen the spreading group's effect compared to a random selection of seeds are (i) a larger dense cluster D' within the graph (left-hand figure), (ii) higher levels of susceptibility of the message, meaning larger values of P0 (middle figure) and (iii) the slower decay in susceptibility (right-hand figure). Note that point (ii) and point (iii) emerge from a rather similar effect, one in which higher levels of susceptibility might result from a higher initial value of P0, as well as a slower decay in retention loss due to smaller values of the decay factor c.

5.2 Model evaluation using the Twitter data

In this section, we utilize the insights from the model to analyze the spread of the Twitter messages related to the NASDAQ-100. We consider several possible scenarios by which

information can spread through a social network. Distinguishing between the scenarios is not always simple, because, given that the real social network structure in Twitter is hidden, the real flow of information is difficult to determine.

The first scenario, called the "naive spreader", corresponds to a case in which an unknown user posts an interesting new message on Twitter and causes a massive information cascade. The user has no desire to create this information cascade. No effort is intentionally invested to increase the spread, but the cascade emerges naturally simply because of the importance of the news and the user's desire to spread it.

The second scenario, called the "guru spreader", is associated with a guru being followed by his/her followers. In this scenario, an initial message tweeted by a well-known user (a guru) is spread by several users (followers), who find this message interesting. An example of such a guru with regard to stock investments is Warren Buffett. In this case, the spread is due to a very connected individual or a highly recognized individual, and the spread is due to his/her degree of connectivity.

The third scenario, on which we focus, is the "spreading group", where a distinguishable group of users intends to spread a message by repeating it frequently within the group. Such a strategy might hide or present the initial spreader. Furthermore, in many cases, unlike the guru, who is a leading figure, this strategy might be used for commercial reasons, such as spreading promotions. Thus, we expect it to contain messages of less economic value because their source is not an expert or a guru. Since in this last scenario the spread is the result of an intentional activity whose goal is solely the spread, it is not a spontaneous act of information cascade as in the two previous cases. Such a message-flow scenario would result in a spread that resembles the spread through the most influential users in the network.

Distinguishing between the scenario of a guru followed by a group of followers and that of a spreading group is not always easy because the Twitter network is hidden from the crawlers. The retweet mechanism hides the real network links, and one cannot always know who follows whom. A guru might create a group of followers who form a cluster with a higher degree of connectivity. The formation of a dense cluster around a financial guru is only a result of his/her financial acumen. To distinguish between groups that are formed around a financial guru and those that are only a method to increase the message spread, we conducted an evaluation of the message's content that we describe at the end of the result sub-section.

5.3 Twitter data and methods of analysis

Of the 1,131,625 unique tweets that included the NASDAQ-100 stock hashtag, we selected 27,071 (2.4 percent of the data) cascades of messages that repeated more than once. We divided these message cascades into three subsets: a subset of tweets with more than 700 repetitions (High group), a subset of tweets that repeated between 401 and 699 times (Mid group) and a subset with tweets repeating between 100 and 400 times (Low group).

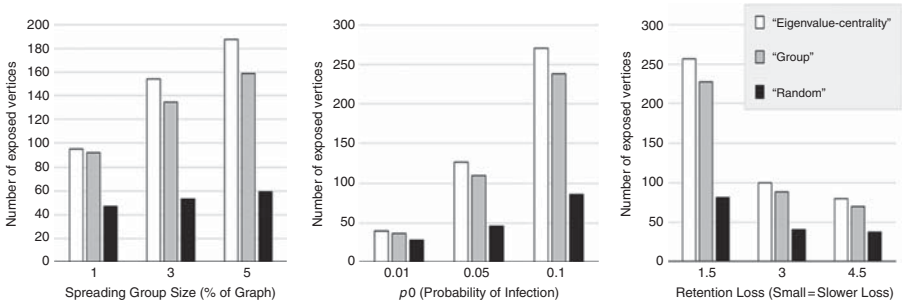

Figure 3. Average number of exposed vertices for seeding group by percent (1 percent, 3 percent, 5 percent) left, by initial retention loss – ($p0$ = 1 percent, 5 percent, 10 percent) center, and by retention loss factor (c = 1.5, 3, 4.5), right

The High group messages contained 79 unique cascades that repeated 2,606 times on average. In this subset, two messages with extremely high levels of repetitions (57,026 and 27,579) were omitted and considered outliers. We removed these two messages since they were an order of magnitude 10 and 5 times higher than the next cascade, but only change 0.03 percent of the entire data set. The next largest cascades were of 5,729, 4,520, 2,427, 1,923, 1,726, and 1,722 repetitions and were kept in the data set. This procedure left us with 77 cascades with an average repeating rate of 1,575 repetitions. The group with the lowest number of repeated messages (Low group) contained 143 unique tweet cascades that repeated 173 times per message on average. To create two distinct groups, we ignored the middle group and considered only the High and Low groups. Given that the number of messages in the two groups was substantially different, we normalized the Low group. To obtain a group of messages in the Low group that was comparable to that in the High group (the High group consisted of 77 messages, while the Low group contained 143 messages), in each run 77 tweets were randomly chosen from the Low group and compared to the High group. As a result, the analysis was conducted with groups of similar sizes, each of which contained 77 unique tweet cascades. The process of randomly choosing users was repeated 20 times to ensure that this random choice did not influence the final results. We compared the messaging activity for users in both the High and Low groups, including a comparison of the rates of reoccurrence of similar users and the distributions of the number of tweets in each group.

5.4 Results

The following analysis addresses the reoccurrence of the tweets and distributions of users in both groups, the High and Low repeating cascades. The main objective of the analysis is to identify differences in tweeting behaviors with respect to the behavior of users involved in the spread of High and Low groups. We assessed the users involved in spreading the messages in these two groups. As indicated above, both groups now contained 77 unique tweet cascades with an average of 2,606 and 173 repetitions (retweets), respectively.

First, for both groups, we identified all the users involved in spreading the cascades of Twitter messages in the High and Low groups. We denoted those by V_h and V_l, respectively. Note that since the same messages can be sent by different users and, alternatively, different messages can be sent by the same users, each of the $\vec{V}_h$ and $\vec{V}_l$ might include overlapping user identifiers in both sets. The number of tweets involved in the spread of messages in each group were $\#\vec{V}_h = 204,442$ and $\#\vec{V}_l = 16,919$.

Our aim was to inspect the characteristics of users that were involved in spreading high/low tweet cascades. To accomplish this, we first randomly chose 20 messages from the High and 20 messages from the Low groups. We then identified all the users that were involved in the spread of each cascade, and randomly selected from each cascade 10 users that were involved in its spread. This resulted in $20 \times 10 = 200$ users that were involved in the spread of messages from V_h and 200 users that were involved in the spread of messages from V_l. Then we collected all of the tweets that were sent by each of these 400 users (200 users from High and 200 from Low) and analyzed the histograms of their tweets, with respect to their number of repetitions. After removing tweets that had no repetitions, we plotted the repetition histograms for tweets that were generated from users in the two groups, as presented in Figure 4. Our goal was to check if the users involved in a single highly repeating cascade can be differentiated in some aspect from users involved in a lowly repeating cascade. Surprisingly, the answer was yes. Figure 4(a) illustrates the histograms of repetition of messages for users in the High group and Figure 4(b) provides the same information for users in the Low group.

In these figures, the cascade length (x-axis) is plotted vs the number of occurrences of a cascade of this length in the entire data set (y-axis). Note that the distributions of retweets for

users in High (Figure 4(a)) vs Low group (Figure 4(b)) are significantly different. Moreover, the distribution of the Low group follows the power law distribution, as expected for an information cascade that does not contain any spreading group scenario, unlike the distribution for the High group, where some users were involved in substantially more tweets.

5.5 Investigating reoccurrence rates of users in High and Low spreading groups

The goal of this section is to describe the procedure used to investigate the early spreaders of a tweet cascade, and the differences in the spreading patterns between High and Low spreading groups[2].

According to the spreading groups model, the users that form the spreading group are more involved in the early stages of the cascade, and their involvement should result in longer cascades. We thus expect longer cascades to have more users belonging to spreading groups at the early stages of the cascade. But can we identify the users belonging to the spreading groups? In our case, we do not know who belongs to the spreading groups and cannot clearly identify them. What we can do is predict if spreading groups exist, by inspecting if some users are over-represented in the early stages of the spread for messages that form longer cascades. Our goal is thus to investigate whether there are differences between the frequency of appearance for early spreaders of messages in the High vs Low cascades. Our assumption is that if spreading groups do exist, and they actually promote the spread, then we need to find more similar users (users from a spreading group) at the early stages of a longer cascade, i.e. High group, as compared to the Low group.

To inspect the above assumption, we performed the following procedure. First, we filtered the data set to create two tables of tweets records, one belonging to the High and the other belonging to the Low repeating cascades. As defined previously in the results section, these two tables, defined as V_h and V_l, included the following fields for all cascades of types High or Low: (1) The tweet's $< \text{text} >$, (2) The tweet's $< \text{user id} >$, and (3) the tweet's $< \text{timestamp} >$. We defined a threshold value m, and collected for all the tweets in V_h and V_l the IDs of the earliest m spreaders for each cascade. We denoted these vectors of m early spreaders by V_h^m and V_l^m (with h and l denoting the High and Low groups). Note that similar ID numbers can appear several times in V_h^m and V_l^m, since these tables include all 77 cascades. Next, we counted how many unique user IDs existed in V_h^m and V_l^m, and denoted these vectors of IDs as $V_h^{m\_u}$ and $V_l^{m\_u}$. Last, we compared the rates of repeating names in both the high and the low vectors by comparing the unique and non-unique vectors. If spreading groups exist in our data, we expect a higher repetition rate of users in the High as compared to the Low group at early stages of the cascade. The rates of reoccurrence of tweets R^m for the High and Low groups were calculated by Equation (3), where the symbol # denotes the size of the vector; the numerator is thus the number of

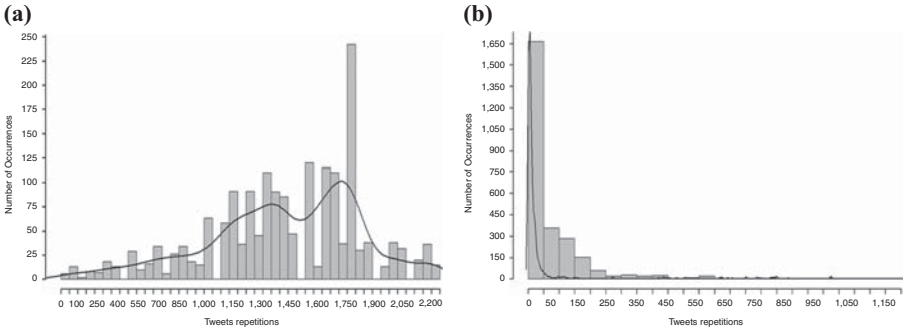

Figure 4.
Histograms of tweet repetitions for the users involved in the High (a) and Low (b) repeated message cascades groups

repeated appearances of users over their first appearance, and the denominator is the total size of names (rows) in the tweet tables. This formula yielded the fraction of repeating names for both groups:

$$R_l^m = \frac{\#V_h^m - \#V_h^{m\_u}}{\#V_h^m} \text{ and } R_h^m = \frac{\#V_l^m - \#V_l^{m\_u}}{\#V_l^m}. \quad (3)$$

To demonstrate the use of this measure, let us consider the following example. Let us assume two cascades of size ten users, as shown in Table I. The first cascade includes the user IDs {100, 101,100,101, …, 101} and the second cascade includes user IDs {100, 101, 102, 103, 100, 101, …, 104}. In the first cascade (Case 1) only two users repeatedly spread the message, while in the second case (Case 2) the message was spread by five users.

In both cases, the total number of messages is $\#V^{10} = 10$, but in case 1 there are only two unique IDs (100 and 101). Thus, the measure of the reoccurrence of the tweets is $R_{case1}^{m=10} = (10-2/10) = 0.8$, while for Case 2 $R_{case2}^{m=10} = (10-5/10) = 0.5$. For a smaller sample, $m = 6$, the values would be $R_{case1}^{m=6} = (6-2/6) = 0.667$, and $R_{case2}^{m=6} = (6-5/6) = 0.166$.

Let us now consider the dependence of rates R_h^m and R_l^m on the number m of the earliest spreaders, determined from our examination of the real Twitter data set, for the High and Low groups. This dependence is presented in Figure 5, while the data and code are available in the GitHub page related to this work.

As the figure shows, for a small value of m, the number of repeating users in both groups is also small and coincides with a low probability of the reoccurrence of the messages. However, as m increases, the difference between the number of repeating users grows. Indeed, for m close to 200 in the group of High messages, this number is nearly twice as high as the repetitions in the Low group. Furthermore, one would assume that messages with a high retweet count would include "fewer" repeating names, not "more" as in the results above, only if no spreading groups exist. This assumption is based on the fact that it is less likely that a random repeating name would appear when more users are involved in the

Table I. Examples of users appearing in two Twitter cascades

T (time)	1	2	3	4	5	6	7	8	9	10
Case 1: ID	100	101	100	101	100	101	100	101	100	101
Case 2: ID	100	101	102	103	104	100	101	102	103	104

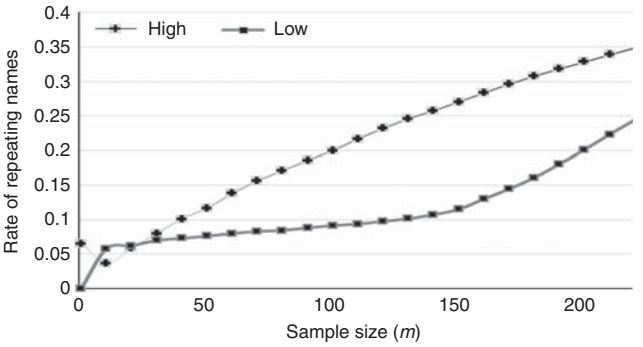

Figure 5. Proportion of repeating names for High and Low repeated tweets groups and sample size m

message's spread. Note, however, that the opposite case is true in the case of real data, as expected for the spreading groups model. While this trend is clear for the early spreaders, as we inspect growing values of m, the rate of similar users in the High group would cease being higher than that in the Low group.

5.6 High-level evaluation of message content vs spread

Since the actual spreading network is hidden in Twitter when using the streaming API crawling process, it is difficult to distinguish between a guru who has many followers and a spreading group. When we investigated a sample of tweets, we found some messages that were spread quite often, containing information that seemed to have little financial value but was spread in large cascades. We suspected that these highly retweeted messages, which did not contain much content, were the result of a spreading group activity. This result led us to investigate the messages' content as a means of distinguishing between the guru scenario and the spreading groups scenario.

From the data set, we selected 12 users who had the highest repetition rate. We investigated these users and found seven users from the High group, which are also known as global financial medial news broadcasters, such as CNBC, ForbesTech, FortuneMagazine, and BloombergTV, and investigated the content of their most retweeted messages. In addition, we identified a second group of five users who, to the best of our knowledge, are not considered global financial media news broadcasters. Examples of such users are accounts with names like "Teacuppiglets," "Philstockworld" and "2morrowknight." In both groups, we examined the number of tweets and retweets. Table II presents the results sorted by the average number of retweets and the total number of retweets.

The average number of retweets for the three large global financial broadcasters, i.e. CNBC, The Wall Street Journal, and ForbesTech, was 42, 125, and 30, respectively. In comparison, the average number of retweets of the messages posted by Philstockworld and Teacuppiglets was 1,291 and 710, respectively.

A high-level comparison of the text content for the top messages sent by the global financial media firms, such as BloombergNews, ForbesTech, and YahooFinance, and the content of the messages sent by the accounts not considered global financial broadcasters demonstrates that tweets by the former often included fundamental information that could be relevant to investors. However, tweets by the latter have less valuable content. For example,

User name	Total number of tweets	Number of retweets	Average number of retweets
Well-known users			
CNBC	113	4,781	42
WSJ	112	13,949	125
ForbesTech	36	1,096	30
CNBCnow	25	736	29
BloombergNews*	23	575	25
FortuneMagazine	57	1,186	21
YahooFinance	87	1,606	18
Less-known users			
Philstockworld	35	45,198	1,291
Teacuppiglets*	44	31,240	710
androsForm	32	3,507	110
ValaAfshar	26	1,174	45
2morrowknight	35	1,073	31

Table II. The number of tweets and retweets for global companies and private users

the message, "RT@teacuppiglets: Bought 6,000 shares of $fb here let's go:)" was spread by 1,051 users, and may (or may not) represent a group of investors following a financial guru. Other examples are the following: "RT@teapartymobile: "$AAPL about to have a huge day! Gapping up already! Weeeeeeeeeeeeeeeeee!!! $MLCG $FB," which was retweeted by 1,726 users, or the message "RT @teacuppiglets: Our private $FB investors group has almost reached 1,500 members. There is power in numbers! Free 2 join it is free https […]" which was retweeted by 1,575 users. These might be considered a spreading group artifact.

5.7 Experiment to evaluate message importance relative to its spread

Our own evaluation of tweet content might be subjective. We thus conducted an experiment aiming to quantitatively evaluate message content vs its spread. We used 100 responders from the Amazon Mechanical Turk service (MT) to review the content of the 20th most retweeted messages from two selected accounts, with a total of 40 messages. We selected from each group in Table II above one Twitter account considered to be a global media broadcaster (i.e. BloombergNews) and another not considered to be a global media broadcaster (i.e. Teacuppiglets). The account BloombergNews was selected because it had the lowest average retweeting rates (thus we expect its content to be the least important). The account Teacuppiglets was selected since its Google search resulted in no relevant links to connect this account to any financial guru or trading company or global media broadcaster.

From the 100 MT responders, we ignored 3 responders due to reliability issues, and were left with 97 responders. Each responder was asked to evaluate 10 questions on the importance of 10 tweets that were published either by BloombergNews or by Teacuppiglets, in terms of his/her own belief of the financial importance of tweet content to a stock market investor. The importance was evaluated on a scale of 1 to 5, with 5 being highly important.

We used 4 different versions of questionnaires[3], with questions shuffled between the four versions to eliminate possible bias due to the position of the question in the questionnaire itself. In each tweet message, the name of the spreader was hidden to prevent any possible bias (negative or positive) due to known media financial broadcasters. The final evaluation of message content is presented in Figure 6, and is the summation histogram of 970 evaluations of 40 different tweets, as performed by 97 users on a 1 to 5 scale.

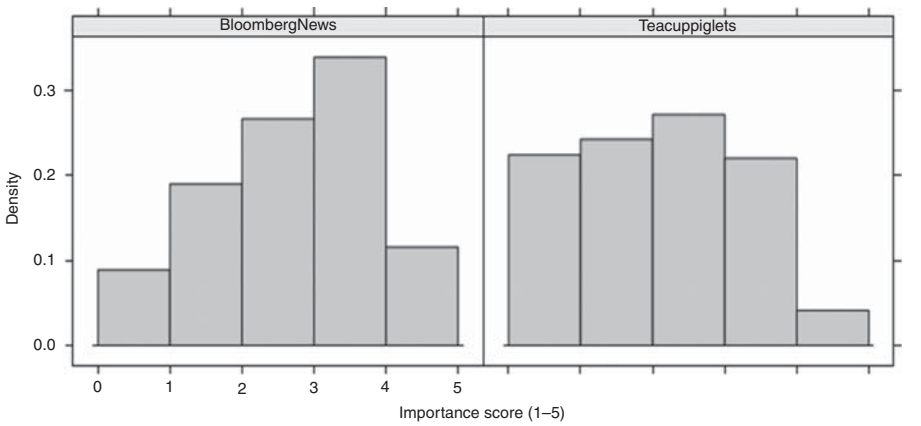

Figure 6.
Histograms of tweet repetitions for the users involved in the High (a) and Low (b) repeated message cascades groups

Notes: 97 distinct reviewers evaluated a total of 970 tweets evaluations, which revealed potentially greater importance for messages that were spread by a global financial media broadcaster

The results obtained also support the assumption that the high degree of spreading (an average of 710 retweets for the unknown source, Teacuppiglets, vs 25 for BloombergNews) does not necessarily relate to the importance of message content, since the messages from Teacuppiglets did not contain content that was more valuable.

6. Discussion

Researchers have documented the use of social bots that mimic human behavior in social networks. Their effect and influence on political, social, and economic aspects has reached such a degree that DARPA, the American Defense Advanced Research Projects Agency, held a 4-week competition with a $4m prize to identify "influence bots" in Twitter (Subrahmanian *et al.*, 2016; Davis *et al.*, 2016). The winning team achieved an accuracy rate of 97.5 percent.

While DARPA's competition winners seem to differentiate very well between social bots and real humans, this is not the case for some naive users nor for the social network sites themselves. For example, Freitas *et al.* (2015) constructed 120 social bots and followed their social development over time. All of these bots had a Twitter profile, were following other users, and generated a random stream of tweets. These bots soon started to be followed by naive users. Moreover, only 31 percent of the bots were detected by Twitter spam services within a period of one month, at the time that the article was published. More interestingly, three of these bots achieved levels of popularity and influence similar to those of the top data scientists at Facebook, even surpassing the principal research members of IBM, as measured by the Klout score (Anger and Kittl, 2011). These results demonstrate the growing influence of social bots and the complexity of their activities, as well as the awareness required for their developing strategies.

Estimations for the percentage of social network accounts that are actually bots (and therefore not human) vary between 9 and 15 percent of the total number of active accounts (Varol *et al.*, 2017). Furthermore, spreading group activity has been shown to be effective not only in stock market-related topics, but also in politics and election power struggles (Echeverria and Zhou, 2017; Ferrara, 2017).

One interesting finding regarding the estimated operating costs and maintenance of spreading groups is that no less than 350,000 bots are synchronized to randomly tweet messages from the Star Wars movie series (Bastos and Mercea, 2019). While these "Star Wars" bots do not necessarily have to be spreading groups, they do represent the ability to automatically create social bots, suggesting very low operating costs for large numbers of bots. Otherwise, it would be hard to imagine that someone manually opens 350,000 Twitter accounts and runs 350,000 bots simply to tweet random Star Wars quotes.

While detecting a single social bot seems to be a task for organizations like DARPA or Facebook, the next level of complexity in bot evolution will include the use of bot spreading groups. These bot networks are intermixed with human users, either naive or non-naive, which makes their identification harder.

Our work demonstrates by simulations the effectiveness of a strategic construction of sub-networks of users to promote the spread of information. With the development of AI technologies such as deep neural networks in text-related tasks (Freitas *et al.*, 2015), as well as the increasing availability of Generative Adversary Networks (Kong *et al.*, 2019), smarter construction of spreading groups might be used to enhance, increase and manipulate the exposure of selected messages on social networks.

While the spreading group model demonstrates a clear advantage of bots operating in synchronized groups, one limitation of our work is that it may be questionable if an account that seems to be a part of a spreading group is a bot or a real human being. This detection difficulty was used to weaken the spreading groups theory, where politicians in the 2019 Israeli elections used a false detection of a bot, which was actually found to be a real person, to claim bots do not exist at all, nor operate in the election (Mail online, 2019; Harkov, 2019).

While currently the detection of spreading groups is based on their activity at early spreading stages, future bots can change their spreading strategy to avoid this method of detection. This might be another limitation.

7. Conclusion

Social networks are an important aspect of modern information spread. Their influence on the social, economic and political aspects of our lives is evident. While a common practice used by commercial organizations to increase their influence and popularity focuses on promoting a message through influencers (e.g. celebrities or opinion leaders), the engagement of these celebrities to promote commercial products (or social agendas) is usually costly and dependent on their personal preferences. This work presents the advantages of a more technical (but probably less ethical) method of achieving the same goal by increasing the influence of a promoted message in a social network using spreading groups.

While the creation of bots and spreading groups is clearly unethical, since it includes a deceptive creation of a fake personality, as long as the legal systems permit this unethical act, this method would probably continue. Initial action toward legal enforcement to prevent fake accounts (and bots) has been made (e.g. by the French president), through restriction on social networks companies during elections (Su, 2018). Nevertheless, these restrictions are only in their early stages.

These spreading groups consist of a dense subset of users in the network, which ignite the spread of a particular message through a ping-pong-like exchange of that message in its initial spreading phase. Through this method, they create an information cascade by passing the message between the members of the group in its early stages, thereby exceeding the natural threshold of the viral process. Thus, the message spreads into larger portions of the network with greater ease.

More and more evidence accumulates for the existence of spreading groups. In some cases, they are referred as BotNets (Hackathon, 2019); nevertheless, BotNets and spreading groups are only synonyms for a similar phenomenon. Although the less sophisticated spreading groups are easy to detect because of their lack of meaningful content, it is only a matter of time until their detection by content will be much harder. Other methods for detecting spreading groups might include more sophisticated cyber methods: e.g., through the IP of their initial account, through the time of account creation, or through the methods proposed in this article.

In the age of social networks, where social connections are a valuable asset, this work has three main implications: first, the awareness, by which we hope to alert the reader to the manipulations that can be (and are) performed by the spreading groups mechanism in order to increase the spread of a message or an agenda; second, the practical aspect, whereby one might want to use the spreading group mechanism, as long as it does not violate any legal restriction, to better spread one's messages; third, the detection, wherein this work only reveals the top of the spreading group iceberg, and we expect much work can be done regarding this in the near future, both by the academic and by the industrial sectors.

Notes

1. https://help.twitter.com/en/using-twitter/retweet-faqs
2. The code and data analysis of this part can be found in the Github account of this work, https://github.com/alonsela2015/OIR-Article
3. Questionnaires of Amazon Mechanical Turk experiment can be found on: http://goo.gl/forms/bKcfl7wELx in http://goo.gl/forms/kJwIMmr3PE in http://goo.gl/forms/roiUtT4LpE and in http://goo.gl/forms/zrzMrHBriy

References

Allcott, H. and Gentzkow, M. (2017), "Social media and fakenews in the 2016 election", *Journal of Economic Perspectives*, Vol. 31 No. 2, pp. 211-236.

Anger, I. and Kittl, C. (2011), "Measuring influence on Twitter", *Proceedings of the 11th International Conference on Knowledge Management and Knowledge Technologies*, p. 31.

Antweiler, W. and Frank, M.Z. (2004), "Is all that talk just noise? The information content of internet stock message boards", *The Journal of Finance*, Vol. 59 No. 3, pp. 1259-1294.

Aral, S., Muchnik, L. and Sundararajan, A. (2013), "Engineering social contagions: optimal network seeding in the presence of homophily", *Network Science*, Vol. 1 No. 2, pp. 125-153.

Barabasi, A.-L. and Albert, R. (1999), "Emergence of scaling in random networks", *Science*, Vol. 286 No. 5439, pp. 509-512.

Bastos, M.T. and Mercea, D. (2019), "The Brexit Botnet and user-generated hyperpartisan news", *Social Science Computer Review*, Vol. 37 No. 1, pp. 38-54.

Bennett, W.L. and Segelberg, A. (2011), "Digital media and the personalization of collective action: social technology and the organization of protests ageins the global economic crisis", *Information, Communication & Society*, Vol. 14 No. 6, pp. 770-799.

Bollobas, B. (2001), *Random Graphs*, Cambridge University Press, Cambridge, MA.

Borgatti, S.P. (2005), "Centrality and network flow", *Social Networks*, Vol. 27 No. 1, pp. 55-71.

Cao, H.H., Coval, J.D. and Hirshleifer, D. (2002), "Sidelined investors, trading-generated news, and security returns", *Reviews of Financial Studies*, Vol. 15 No. 2, pp. 615-648.

Chierichetti, F., Kleinberg, J. and Panconesi, A. (2014), "How to schedule a cascade in an arbitrary graph", *SIAM Journal on Computing*, Vol. 43 No. 6, pp. 1906-1920.

Davis, C.A., Varol, O., Ferrara, E., Flammini, A. and Menczer, F. (2016), "BotOrNot: a system to evaluate social bots", *Proceedings of the 25th International Conference Companion on World Wide Web*, *International World Wide Web Conferences Steering Committee*, pp. 273-274.

DeMarzo, P.M., Vayanos, D. and Zwiebel, J. (2003), "Persuation bias, social influence, and unidimensional opinions", *The Quarterly Journal of Economics*, Vol. 118 No. 3, pp. 909-968.

Echeverria, J. and Zhou, S. (2017), "Discovery, retrieval, and analysis of the Star Wars´Botnet in Twitter", *Proceedings of the 2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2017*, pp. 1-8.

Ferrara, E. (2017), "Disinformation and social bot operations in the run up to the 2017 French presidential election", *First Monday*, Vol. 22 No. 8.

Freitas, C., Benevenuto, F., Ghosh, S. and Veloso, A. (2015), "Reverse engineering socialbot infiltration strategies in twitter", *Proceedings of the 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2015*, pp. 25-32.

Funk, S., Gilad, E., Watkins, C. and Jansen, V.A. (2009), "The spread of awareness and its impact on epidemic outbreaks", *Proceedings of the National Academy of Sciences*, Vol. 106 No. 16, pp. 6872-6877.

Goldenberg, D., Sela, A. and Shmueli, E. (2018), "Timing matters: influence maximization in Social networks through scheduled seeding", *IEEE Transactions on Computational Social Systems*, Vol. 5 No. 3, pp. 621-638.

Gorman, G.E. (2014), "The key is trust': communicating meaningfully in an online environment", *Online Information Review*, Vol. 38 No. 1, p. 38.

Hackathon, T.B.-H. (2019), "The center for cyber, law and policy", February 7-9, available at: https://cri.hevra.haifa.ac.il/events/2019/the-bot-hunting-hackathon (accessed February 24, 2019).

Harkov, L. (2019), "Jerusalem post", available at: www.jpost.com/Israel-Elections/From-Begins-tchach-tchachim-to-Bibis-bots-Analysis-585657 (accessed May 6, 2019).

Harris, M. and Raviv, A. (1993), "Differences of opinion make a horse race", *Review of Financial Studies*, Vol. 6 No. 3, pp. 473-506.

Harvey, K. (2013), *Encyclopedia of Social Media and Politics*, Sage Publications.

Hirshleifer, D. and Teoh, S.H. (2003), "Limited attention, information disclosure, and financial reporting", *Journal of Accounting and Economics*, Vol. 36 No. 1, pp. 337-386.

Howard, P.N., Duffy, A., Freelon, D., Hussain, M.M., Mari, W. and Mazaid, M. (2011), "Opening closed regimes: what is the role of social media during the Arab Spring?", Social Science Research Network, available at: http://ssrn.com/abstract=2595096 (accessed February 4, 2016).

Kempe, D., Kleinberg, J. and Tardos, E. (2003), "Maximizing the spread of influence through a social network", *Proceedings of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 137-146.

Kim, J. and Hastak, M. (2018), "Social network analysis", *International Journal of Information Management: The Journal for Information Professionals*, Vol. 38 No. 1, pp. 86-96.

Kong, X., Li, B., Neubig, G., Hovy, E. and Yang, Y. (2019), "An adversarial approach to high-quality, sentiment-controlled neural dialogue generation", arXiv preprint, arXiv: 1901.07129.

Leskovec, J., Backstrom, L. and Kleinberg, J. (2009), "Meme-tracking and the dynamics of the news cycle", *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 497-506.

Mail online (2019), "Israel seeks to beat election cyber bots", GMT, February 1, available at: www.dailymail.co.uk/wires/afp/article-6656155/Israel-seeks-beat-election-cyber-bots.html (accessed May 26, 2019).

Michalski, R., Kajdanowicz, T., Br'odka, P. and Kazienko, P. (2014), "Seed selection for spread of influence in social networks: temporal vs. static approach", *New Generation Computing*, Vol. 32 Nos 3-4, pp. 213-235.

Newman, M. (2010), *Networks: An Introduction*, Oxford University Press, Oxford.

Newman, M.E. (2005), "A measure of betweenness centrality based on random walks", *Social Networks*, Vol. 27 No. 1, pp. 39-54.

Newman, N., Fletcher, R., Kalogeropoulos, A., Levy, D. and Nielsen, R.K. (2017), "Reuters Institute digital news report 2017".

Sela, A., Goldenberg, D., Ben-Gal, I. and Shmueli, E. (2017), "Latent viral marketing, concepts and control methods", arXiv preprint, arXiv: 1704.01775.

Sela, A., Goldenberg, D., Ben-Gal, I. and Shmueli, E. (2018), "Active viral marketing: incorporating continuous active seeding efforts into the diffusion model", *Expert Systems with Applications*, Vol. 107, pp. 45-60.

Sela, A., Shekhtman, L., Havlin, S. and Ben-Gal, I. (2016), "Comparing the diversityof information by word-of-mouth vs. web spread", *Europhysics Letters*, Vol. 114, pp. 58003.

Sela, A., Shmueli, E., Goldenberg, D. and Ben-Gal, I. (2016), "Why spending more might get you less, dynamic selection of influencers in social networks", *2016 IEEE International Conference on the Science of Electrical Engineering (ICSEE)*, pp. 1-4.

Shakarian, P., Eyre, S. and Paulo, D. (2013), "A scalable heuristic for viral marketing under the tipping model", *Social Network Analysis and Mining*, Vol. 3 No. 4, pp. 1225-1248.

Su, J.B. (2018), "France to impose restrictions on Facebook, Twitter in fight against fake news during elections", *Forbes*, January 9, available at: www.forbes.com/sites/jeanbaptiste/2018/01/09/france-could-ban-facebook-twitter-in-fight-against-fake-news-during-elections/#7f82522749da (accessed May 26, 2019).

Subrahmanian, V.S., Azaria, A., Durst, S., Kagan, V., Galstyan, A., Lerman, K., Zhu, L., Ferrara, E., Flammini, A. and Menczer, F. (2016), "The DARPA Twitter Bot Challenge", *Computer*, Vol. 49 No. 6, pp. 38-46.

Teng, S., Wei Khong, K., Wei Goh, W. and Yee Loong Chong, A. (2014), "Examining the antecedents of persuasive eWOM messages in social media", *Online Information Review*, Vol. 38 No. 6, pp. 746-768.

US Securities and Exchange Commission (1998), "SEC charges 44 stock promoters in first internet securities fraud sweep", US Securities and Exchange Commission, October 28, available at: www.sec.gov/news/headlines/netfraud.htm

Varol, O., Ferrara, E., Davis, C.A., Menczer, F. and Flammini, A. (2017), "Online human-bot interactions: detection, estimation, and characterization", *Eleventh International AAAI Conference on Web and Social Media*.

Weng, L., Flammini, A., Vespignani, A. and Menczer, F. (2012), "Competition among memes in a world with limited attention", *Scientific Reports*, Vol. 2, p. 335.

Wilensky, U. (1999), "NetLogo", available at: https://ccl.northwestern.edu/netlogo (accessed February 6, 2016).

Wilensky, U. and Rand, W. (2015), *An Introduction to Agent-Based Modeling: Modeling Natural, Social, and Engineered Complex Systems with NetLogo*, MIT Press, Cambridge, MA.

Wu, F., Huberman, B.A., Adamic, L.A. and Tyler, J.R. (2004), "Information flow in social groups", *Physica A: Statistical Mechanics and its Applications*, Vol. 337 Nos 1-2, pp. 327-335.

Zhou, C., Zhao, Q. and Lu, W. (2015), "Impact of repeated exposures on information spreading in social networks", *PloS One*, Vol. 10 No. 10, p. e0140556.

Corresponding author
Alon Sela is the corresponding author and can be contacted at: alonse@ariel.ac.il