

Understanding Popularity, Reputation, and Social Influence in the Twitter Society

David Garcia, Pavlin Mavrodiev, Daniele Casati, and Frank Schweitzer

The pervasive presence of online media in our society has transferred a significant part of political deliberation to online forums and social networking sites. This article examines popularity, reputation, and social influence on Twitter using large-scale digital traces from 2009 to 2016. We process network information on more than 40 million users, calculating new global measures of reputation that build on the D-core decomposition and the bow-tie structure of the Twitter follower network. We integrate our measurements of popularity, reputation, and social influence to evaluate what keeps users active, what makes them more popular, and what determines their influence. We find that there is a range of values in which the risk of a user becoming inactive grows with popularity and reputation. Popularity in Twitter resembles a proportional growth process that is faster in its strongly connected component, and that can be accelerated by reputation when users are already popular. We find that social influence on Twitter is mainly related to popularity rather than reputation, but that this growth of influence with popularity is sublinear. The explanatory and predictive power of our method shows that global network metrics are better predictors of inactivity and social influence, calling for analyses that go beyond local metrics like the number of followers.

KEY WORDS: social network analysis, social influence, online social networks, popularity, reputation

Introduction

The pervasive presence of online media in our society has transferred a significant part of political deliberation to online forums and social networking sites. The picture of a café where citizens discuss and exchange ideas is being replaced by the interface of an online network in which citizens communicate and discuss with almost anyone from almost anywhere (Castells, 2011). The transformative role of these technologies can fundamentally change traditional political processes. An illustrative example is how collective action (Olson, 1965) evolves into the process of connective action (Bennett & Segerberg, 2012), in which citizens deliberate and self-organize in a decentralized way by using information technologies.

The new level of connectivity of our online society raises important questions about the roles that citizens have in the political process, when they become *users*

of online social media (boyd & Crawford, 2012). Understanding important phenomena, such as social influence, social forces, and digital divides, requires the analysis of the evolution of very large social systems, which traditionally has been a challenging task in the Social Sciences. Longitudinal studies so far were only able to monitor small amounts of individuals over time, and pose strong limitations to our knowledge of fundamental social dynamics. Understanding popularity and social influence, or how individuals recognize their peers and adopt their behavior is particularly difficult, as a valid approach requires the observation of natural behavior of many individuals without interference (MacPhee, 1963; Webb, Campbell, Schwartz, & Sechrest, 1966). One solution lies in the vast amount of data produced by our society, which allows us to monitor human behavior at unprecedented sizes, resolutions, and time scales. The field of Computational Social Science seeks to address questions about social behavior with large-scale data sets of digital traces (Lazer et al., 2009), testing hypotheses drawn from theories from the Social Sciences (González-Bailón, 2013).

We present a study of popularity, reputation, and social influence in the Twitter online social network, including more than 40 million users in two snapshots seven years apart. Twitter is a popular online social network in which users establish directional follower links and publicly communicate short messages across them. The extreme popularity of Twitter and its relevance for Political Science (Aragón, Kappler, Kaltenbrunner, Laniado, & Volkovich, 2013; Barberá, 2015; Conover, Gonçalves, Flammini, & Menczer, 2012; Lietz, Wagner, Bleier, & Strohmaier, 2014), makes it an ideal case to study social influence processes on the limits of size and temporal scales. Popularity in the Twitter social network is often quantified as the amount of followers of a user. That implies, it does not matter why some user follows you, or how important he or she is, your popularity only measures the size of your audience. Reputation, on the other hand is a more complicated concept associated with centrality (Friedkin, 1991). Being followed by a highly reputed user has a stronger effect in one's reputation than being followed by someone with low reputation. Thus, the simple amount of followers does not capture the recursive nature of reputation.

Popularity can be compared to reputation using geometric terms. Popular users are at the center of network stars and have high in-degree, while reputable users have high positions in a pyramid, in which they might not be followed by many but their followers are also high in the pyramid. In this article, we aim to unravel the difference between popularity and reputation on the process of social influence (Riquelme & González-Cantergiani, 2016), that is, how information spreads in the Twitter network. We explore the functioning of Twitter combining psychological and economical perspectives to formulate hypotheses, testing them against large-scale data through statistical analyses and predictive methods.

Background and Research Questions

One of the main psychological questions of Twitter use is about user motivation. What are the reasons that keep users active on Twitter and how do

they depend on the attention of others? Empirical studies of user motivation show that it grows with popularity and reputation, but the rate of growth of motivation is sublinear and the risk of leaving the network decreases in the same way (Toubia & Stephen, 2013). This trend of motivation was taken as an initial approximation to study the social resilience of online communities (Garcia, Mavrodiev, & Schweitzer, 2013), but empirical evidence at a large scale is needed to test it and find its limitations. As suggested by Garcia et al. (2013), an alternative hypothesis arises from the phenomenon of information overload in computer-mediated communication (Hiltz & Turoff, 1985). The utility of belonging an online social network can decrease with popularity, in particular, when very active and popular Twitter users are unable to process all the information they receive (Rodriguez, Gummadi, & Schoelkopf, 2014), or when they perceive risks associated to loss of privacy. We formulate this scenario through the *non-monotonic motivation hypothesis*, by which the probability to become inactive would have a U-shape function with respect to popularity and reputation.

Previous works on social resilience suggest that centrality measures are better predictors of user inactivity than degree alone (Garcia et al., 2013; Yu et al., 2016). When users become inactive, they might produce cascades of other users becoming inactive, leading to effects that go beyond their local neighborhood. Based on this, we test the *cascading inactivity hypothesis*, which states that reputation, a centrality measure, is a better predictor of a user becoming inactive than popularity. Further network aspects can influence the decision to become inactive or not. Agent-based models suggest that nodes that belong to cycles are in more stable subsets of the network (Jain & Krishna, 2001). Inspired by this computational model, we test the *core activity hypothesis*, that is, users that belong to the Strongly Connected Component of Twitter (as explained more in detail below) are less likely to become inactive.

The psychological perspective on the functioning of Twitter allows to explain a variety of phenomena relevant for policy, including social influence. Social influence is considered in social impact theory as a manifestation of social impact that leads to a change of behavior after social interaction (Latané, 1981). Social impact theory allows us to formulate testable hypotheses about how social influence depends on group size and source strength. A hypothesis drawn from this theory is the *division of impact*: The larger the addressed group, the weaker the social impact exercised on each person. With respect to Twitter, this translates as a lower influence rate for users with more followers, leading to total influence scaling sublinearly with popularity.

Social impact theory postulates that the extent of impact will increase with source strength, which can be seen as a manifestation of status or reputation (Latané, 1981). If reputation is the driving force of Twitter users, the *reputation strength hypothesis* will hold, that is, reputation is a stronger indicator of social influence than popularity. On the contrary, it is also possible that popularity drives the total amount of attention in Twitter, and thus the amount of followers of a user would be a better indicator of source strength. Additionally, other characterizations of network positions, like the bow-tie structure, can encode

aspects of node influence. For example, companies in the driving group of ownership bow-tie structures hold a strong control over the rest of the network (Vitali, Glattfelder, & Battiston, 2011). This motivates our *bow-tie influence hypothesis*: users in the Out group of the bow-tie decomposition of the Twitter follower network are more influential.

The economic aspects of Twitter are driven by the attention economy in which the scarce resource is attention instead of information (Huberman & Wu, 2008). From an economic perspective, the variable of interest is popularity in the sense of measuring total attention rather than its quality or final effect. Popularity in online media can be subject to capital accumulation phenomena, like the winner-takes-all effect. Following this rationale, we formulate the *popularity growth hypothesis*: the popularity of a user is associated with larger gains of popularity over time.

Furthermore, competition for popularity can have various degrees of strength with respect to reputation. If competition is strong, reputation can decrease the growth of popularity due to the fact that reputable users bring the attention of other reputable users, who are in a stronger position to compete for the attention of the rest of the population. As an alternative, we formulate the *weak competition hypothesis*: reputation has a positive effect on the growth of popularity. This is possible in interaction with popularity, such that reputation can be a complementary resource but not a replacement for popularity. This kind of competition scenarios can be observed in rumor spreading simulations (Borge-Holthoefer & Moreno, 2012) in which the centrality of a user can hinder the spreading of information from less connected nodes.

The economic perspective raises the question of inequality in user attention. Competing for attention often leads to the analysis of the distribution of amount of followers with respect to highly unequal Pareto distributions, or power-laws. Often the *power-law popularity hypothesis* is tested in an uncritical way (Stumpf & Porter, 2012) and a simple linear regression over a log-log distribution plot has been used to state claims about preferential attachment or edge copying. Substantial developments in statistical methods show the problems of that approach (Alstott, Bullmore, & Plenz, 2014; Clauset, Shalizi, & Newman, 2009). Preferential attachment and edge copying, as multiplicative growth mechanisms can also lead to log-normal distributions that display high inequality, but do not have scaling properties (Mitzenmacher, 2004). We aim at properly testing the power-law hypothesis versus the alternative log-normal explanation, evaluating if scaling inequalities prevail over time.

We summarize the nine hypotheses tested in this article in Table 1, sorting them in the order in which they appear in the results section. To test them, we process the full Twitter social network data of 2009 (Kwak, Lee, Park, & Moon, 2010), measuring the popularity and reputation of more than 40 million Twitter users. We trace further in time a large subset of those users, analyzing later information in 2016. Using this later data, we quantify the social influence of each user through information spreading and measure their current popularity and state of activity. First, we analyze the role of reputation and popularity in user motivation, testing if popularity and reputation can have non linear effects in

Table 1. List of the Hypotheses Tested in This Article

Hypotheses on Activity	
Nonmonotonic motivation	The probability of inactivity shows a negative trend with popularity and reputation, but increases for highly reputable and popular users.
Cascading inactivity	Reputation is a better predictor of inactivity than popularity.
Core activity	On average, users in the SCC have lower chances of becoming inactive.
Hypotheses on Popularity	
Power-law popularity	The popularity of active users follows a power-law distribution.
Popularity growth	Popular users tend to gain more popularity over time.
Weak competition	Reputation has a positive effect in the growth of popularity.
Hypotheses on Social Influence	
Division of impact	Social influence grows sublinearly with popularity.
Bow-tie influence	Active users in Out group of the bow-tie are more influential.
Reputation strength	Reputation is a stronger indicator of social influence than popularity.

activity. Second, we analyze the changes in popularity of users, evaluating whether reputation helps or hinders popularity growth, and the level of inequality in popularity. And third, we test if social influence grows sublinearly with popularity and analyze how social influence depends on reputation and popularity, as measured very early in the age of Twitter.

Data and Methods

We study the relationship between popularity, reputation, inactivity, and social influence in Twitter, one of the most widely used and analyzed online social networks (Tufekci, 2014). Twitter is a microblogging social network, meaning that users produce very short posts called *tweets* of up to 140 characters. Users *follow* each other by creating links as subscriptions to the tweets produced by other users. These follower relationships do not need to be reciprocal, in fact only 22 percent of the connections between pairs of users are reciprocated (Kwak et al., 2010). A user can *retweet* a tweet produced by another user, copying that tweet and spreading it to their followers. The asymmetry of information flow through follower links and the power of Twitter to spread information make it a key online social network for celebrities, politicians, news media, opinion makers, and online marketing campaigns.

We study the Twitter follower network, in which nodes are users and directed links connect user i to user j , if and only if user i follows user j . In this network, the in-degree of a node is the amount of followers of a user i , and its out-degree is the amount of users followed by user i . Our analysis leverages a Twitter data set from previous research (Kwak et al., 2010), combining it with new data sets that we produce to track large-scale, longitudinal changes of user properties, and social influence in Twitter. In the following, we summarize the

data sets that we analyze, followed by an explanation of the network and activity metrics that we calculate from those data sets.

Data Sets

Our study starts with the full 2009 network, which includes the public follower lists of all Twitter users in June 2009 (Kwak et al., 2010). The full 2009 network includes 1.47 billion follower links between more than 41 million users, which we take as the user base to analyze in our study. While in earlier times the data set included information on user profiles and the tweets they produced, at the time of writing only the follower network is available, in line with the current Terms of Service of Twitter.¹ This limitation does not affect our network analysis and the public availability of the network data set allows independent replications of our study.

Building on the 2009 user base, we produce a 2016 user status data set in which we look up the status of each user included in the full 2009 network. Using the Twitter API iteratively, we scan the full list of users in the full 2009 network to get updated status data in July 2016. This way, we record the amount of followers in July 2016 and the date of the *last tweet* produced by the user, to have an estimate of popularity and state of activity or inactivity of each user. In total, the API returned data on 35,868,457 users, the missing user profiles either being deleted or set as private.²

Measuring social influence requires data on the activity of users, which is a much more limited data resource in the Twitter API. To study social influence, we focus on the set of active users in 2016, retrieving their latest tweets up to a maximum of 200 tweets in December 2016. Focusing on users that got at least one retweet, we construct the 2016 social influence data set, which includes 351,037,535 tweets from 3,666,717 users. For each of these tweets, we count with the amount of retweets that they got, a quantity that we use to measure social influence as explained below.

Measuring Reputation in Twitter

We measure reputation through the D-core decomposition (Giatsidis, Thilikos, & Vazirgiannis, 2013) of the full 2009 network. The D-core decomposition of a directed network computes a set of (k,l) -cores, defined as maximally large subgraphs of the network in which each node has at least an in-degree of k and an out-degree of l . These (k,l) -cores are identified in a pruning algorithm that removes nodes from cores depending on their in- and out-degree, mapping each user to a set of (k,l) -cores, as explained more in detail in Giatsidis et al. (2013).

The left panel of Figure 1 shows the amount of nodes in each (k,l) -core for the full 2009 network. We observe that most of users belong only to (k,l) -cores with low k and l , while the more central cores with high k and l are smaller in size. The green line shows the bisector that divides the plot in two triangles with the same amount of nodes. It is notable that the bisector is very close to the line of slope 1, with a small higher density toward k than toward l . This indicates a bit larger expected value associated to in-degree (k) than to out-degree (l).

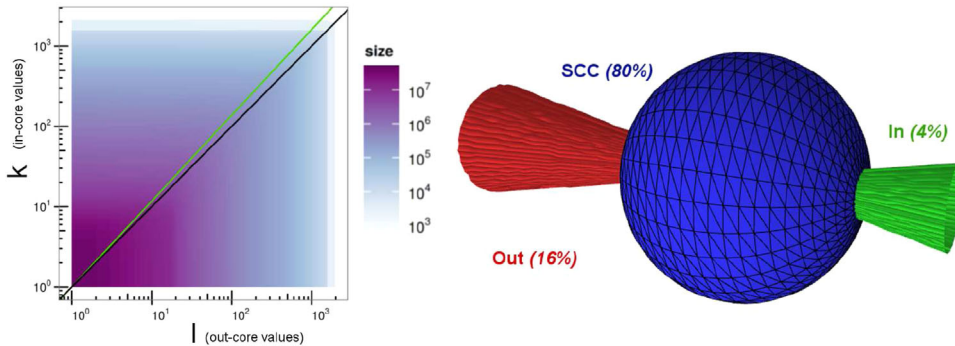


Figure 1. Left: Distribution of (k,l) values of the D-core decomposition of Twitter. All three axes are shown on logarithmic scales, the green line shows the bisector that divides the distribution in two triangles of the same density and the black is the diagonal. Right: Bow-tie structure of Twitter with groups of volume proportional to the amount of users they contain. In blue the SCC is shown (33.5 M users), in red the Out group (6.4 M), and in green the In group (1.5 M). Not shown are nodes not connected to the SCC (167 K, less than 1 percent of the total).

By applying the D-core decomposition, we measure reputation in a global sense by capturing the recursive property of reputation as a centrality measure: Users that belong to (k,l) -cores with high k are followed by users that also belong to (k,l) -cores with high k . While other network centrality measures can capture similar properties, such as HITS (Kleinberg, 1999) or betweenness centrality (Freeman, 1977), the D-core decomposition has two advantages. First, the computation of the D-core decomposition can be run iteratively and scales linearly with the amount of nodes and edges, avoiding costly matrix operations of other measures. It would require extremely large computational power to compute betweenness over the more than 40 million nodes in the full 2009 network. Second, the D-core decomposition is an extension of the k -core decomposition of an undirected network (Seidman, 1983). The k -core decomposition has been shown to be related to spreading processes in real networks (Kitsak et al., 2010), including the spreading of political movements in Twitter (Alvarez, Garcia, Moreno, & Schweitzer, 2015; González-Bailón, Borge-Holthoefer, Rivero, & Moreno, 2011).

Our second approach to reputation focuses on detecting user positions based on the analysis of the *bow-tie* structure of the full 2009 network. Directed networks can be divided in strongly connected components, which are maximal sets of nodes in which every two nodes in the set can be reached through a directed path. A bow-tie structure is characterized by four groups of users: the *SCC group* formed by the largest Strongly Connected Component of the network, the *In group* formed by the users that are connected through a directed path to the SCC but that cannot be reached back from the SCC, the *Out group* with the opposite property, that is, users in the Out group can be reached from the SCC but the SCC cannot be reached back from them, and the rest of users without connections to the SCC. The most popular example of a bow-tie structure is the one present on the Web (Broder et al., 2000), including roughly a fourth of the websites in its SCC.

The right panel of Figure 1 shows the bow-tie structure of the full 2009 network. It shows an example of a giant connected component, with 80 percent of the users in the SCC. Interestingly, the Out group is much larger than the In group, and the rest of users not in any of these three groups is less than 1 percent. We will use this classification in groups as another take on reputation, paying special attention on whether belonging to one group or another is related to different levels of social influence or tendencies to become inactive.

User-Level Metrics

1. Popularity in 2009 and 2016: We measure popularity as the in-degree, that is the amount of followers a user has in Twitter. Similar to the concept of name generators, the amount of followers of a user resembles the amount of times a user is named as being linked to other users, when asking for a question of personal relevance (Niven, Garcia, van der Löwe, Holman, & Mansell, 2015). We count popularity from data sets at two different times, one from the full network in 2009, and a second one from the user status data set in 2016.
2. Reputation as in-coreness in 2009: The D-core decomposition offers a computable basis to extract a reputation measure. We quantify the reputation of a user as the in-coreness of a node, which is the maximum value of k among the (k,l) -cores that the node belongs to. This way, the in-coreness measures reputation as a recursive centrality node property (Friedkin, 1991) related to social spreading processes. Having a high in-coreness is a sufficient condition for having a high in-degree, but not the opposite. Thus, in our measures, reputation is a lower boundary for popularity, but being popular does not imply having a high reputation.
3. Bow-tie group in 2009: Our second metric related to reputation focuses on the importance of the SCC rather than on a scale of reputation. We mark each user as the bow-tie group they belong in the full 2009 network, that is, SCC, In, Out, or Rest. Through this metric, we classify the network in its major core versus a periphery of marginalized users who would not benefit from the recursivity of reputation in the SCC.
4. Inactivity by 2016: Using the timestamp of the last tweet of a user, we classify them as active or inactive. As explained more in detail in the following section, we mark a user as active, if they posted a tweet since April 1, 2016, approximately three months before the data retrieval date in 2016.
5. Social influence as average retweet rate in 2016: We quantify the social influence of a user as the average of the amount of retweets received by the tweets of each user that we captured in the 2016 social influence data set. This metric, inspired in Twitter social influence research (Abisheva, Garimella, Garcia, & Weber, 2014; Guille, Hacid, Favre, & Zighed, 2013), captures the power of a user to spread information across Twitter as a manifestation of immediate social influence, that is, the activity of other users in terms of retweets.

We perform our analysis to evaluate if popularity, reputation, and bow-tie groups measured in 2009 can predict inactivity, popularity, and social influence in 2016. We create user vectors including the relevant metrics presented above, and perform a series of statistical analyses explained below. Online interactive visualizations of the results of this article can be browsed at <https://www.sg.ethz.ch/research/online-social-networks/ipp-twitter-paper>. Codes and data to replicate the results are available at https://github.com/dgarcia-eu/IPP_Twitter_2016.

Inactivity Analysis

We use the last tweet of users as a way to estimate if they are still active or inactive in July 2016. Before calculating descriptive statistics, we randomly exclude 10 percent of the users and save them in a leave-out sample for validation, to evaluate the generalizability and predictive power of our results beyond statistical inference. Figure 2 shows the amount of last tweet timestamps per week on the 2016 data set. A peak can be observed right before the data retrieval date in 2009, illustrating that some users captured in the 2009 data set became inactive around that date. The peak toward the end of Figure 2 is an artifact generated by the users that are still active, as the timestamp of their last tweet is very recent. For our statistical analysis on the detection of inactivity, we focus on all the users that produced their last tweet since July 2009, to avoid oversampling short-lived accounts that were captured only because data was retrieved on that particular date. That leaves us with 15,792,514 users, which we classify as active, if their last tweet was after April 1, 2016, and inactive otherwise. A total of 5,621,145 users were found to remain active (35.5 percent).

We address the questions of whether not belonging to the SCC is associated with inactivity tendencies, whether reputation is a better predictor for inactivity than popularity, and whether the dependence of inactivity on these two is non-monotonic and shows signs of overloads. We first perform nonparametric analyses of inactivity

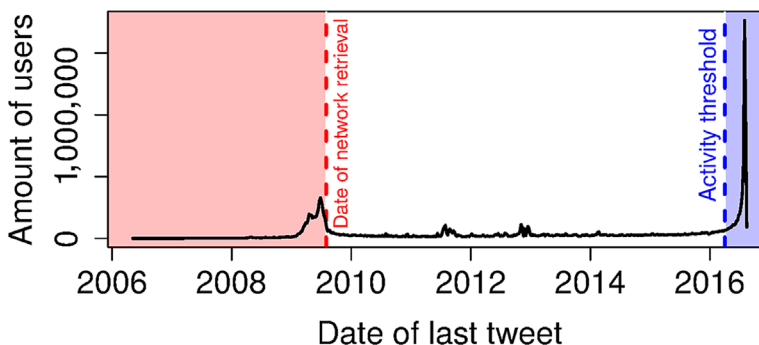


Figure 2. Histogram of the Date of the Last Tweet of Users as of July 2016 (Week Resolution).
Notes: The red vertical line shows the date of the original data retrieval: July 31, 2009. We exclude from the inactivity analysis all the users with a last tweet before that date. The blue vertical line shows the date three months before our analysis, April 1, 2016. Users with a last tweet before that date are considered inactive.

tendencies based on reputation and popularity, to then formalize and fit logistic regression models that we evaluate in a single run over the leave-out sample.

In line with the core activity hypothesis, we find that users in the SCC are substantially less likely to become inactive than those outside it. Two-tailed tests of equality of proportions (Wilson, 1927), between the ratios of inactive users in each bow-tie group show that the probability of inactivity in the SCC is lower than in the In group ($\chi^2=3,887$, $p < 10^{-15}$), the Out group ($\chi^2=53,705$, $p < 10^{-15}$), and the rest ($\chi^2=1,193.8$, $p < 10^{-15}$). These results persist when analyzing only users with both in- and out-degree larger than zero, as explained more in detail in supporting Appendix SI. The fact that users in the SCC are less likely to become inactive suggests that the SCC is a special position in the network, as motivated by previous work on ownership and network dynamics (Jain & Krishna, 2001; Vitali et al., 2011).

Iterating over the popularity and reputation of users, we calculated the ratio of inactive users over 1,000 bootstrap samples. The left panel of Figure 3 shows the estimate of inactivity and its 95 percent bootstrapped confidence interval versus popularity and reputation in 2009. It can be appreciated that the probability of becoming inactive decreases with both reputation and popularity, when these are low, but between a value of 100 and 1,000 they start to increase. After that increase, the ratio of inactive users decreases moderately with reputation, but much stronger with popularity. As such, one can observe that the tail of popularity is associated with lower tendency to become inactive.

The nonlinear shape of the ratio of inactive users versus reputation and popularity suggests the existence of three local optima: a maximum at a very low value, a minimal point between 100 and 1,000, and local maximal point further in the scale. Before modeling, we apply the log-modulus transformation

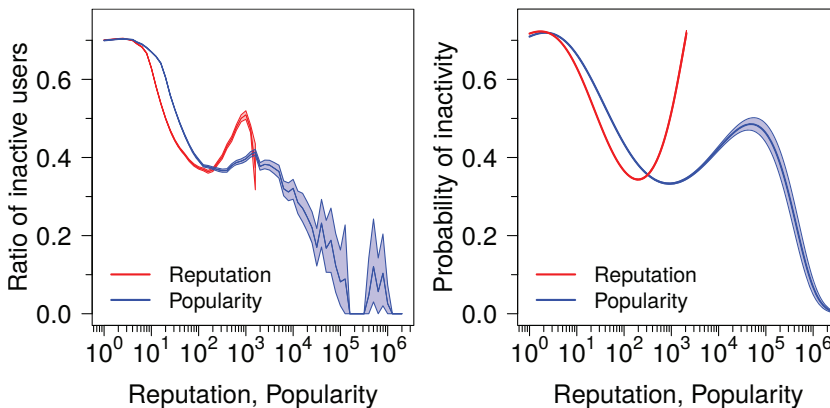


Figure 3. Left: Ratio of inactive users versus reputation and popularity, the shaded area shows 95 percent bootstrapping confidence intervals. Both functions suggest the existence of three local optima in a decreasing but nonlinear trend. Right: Predicted probability of inactivity versus reputation and popularity in their respective models. Shaded areas show the standard error around the prediction. Nonlinear effects are similar to those on the left, with a U-shape for reputation and an additional decrease for very high popularity.

($\ln(x) = \log(x + 1)$) to reputation (in-coreness) and popularity (amount of followers), calculating the values of the random variables I and F , respectively. We model this shape as a four-order polynomial over I and F :

$$\begin{aligned} \text{logit}(P(\text{inactive})) &= b_0 + b_1 \cdot I + b_2 \cdot I^2 + b_3 \cdot I^3 + b_4 \cdot I^4 + \epsilon_I \\ \text{logit}(P(\text{inactive})) &= c_0 + c_1 \cdot F + c_2 \cdot F^2 + c_3 \cdot F^3 + c_4 \cdot F^4 + \epsilon_F \end{aligned} \quad (1)$$

The logit function computes the logarithmic odds ratio of becoming inactive, modeling this as a linear combination of the terms in a fourth order polynomial depending on the log in-coreness I as reputation in the first model, and on the logarithm of the amount of followers F as popularity in the second one. We fit the models using weakly informative priors in the `bayesglm` R package (Gelman, Jakulin, Pittau, & Su, 2008), and present the results in detail in Appendix SII. For both models, all coefficients up to the fourth order are significant, supporting that the choice of a polynomial model was appropriate. The sign of b_4 and c_4 , the fourth order terms is negative, indicating the general asymptotically decreasing tendency of inactivity with reputation and popularity.

Interpreting the detail about the shape of the polynomials is complicated by just inspecting the coefficient estimates. For that reason, we show the predicted values of both models over the range of values of reputation and popularity on the right panel of Figure 3. The predicted probability of inactivity for both reputation and popularity shows a behavior following the nonmonotonic motivation hypothesis. In particular, the reputation function has an increasing tail of the probability of inactivity, and the decreasing tendency of the fourth order term is only present as softening the tail. On the popularity model using followers as predicting variable, the increase in probability is also present, but the decreasing tail toward very high values of amount of followers is present. This indicates that an increase in reputation and popularity from low values is associated with a lower probability of becoming inactive, but there is also a range of values of reputation and popularity in which the opposite effect holds. The difference between model results lies on the decreasing tail, while our observations on the left panel of Figure 3 shows it for both, for the case of reputation the amount of evidence of a decrease at the tail is not sufficient, while for popularity it is clear that very popular users are very unlikely to become inactive.

The reputation model has a better predicting power than the popularity model in terms of the Bayesian information criterion, as shown in Appendix SII. To further test this increased information in the in-coreness, we computed the area under the curve (AUC) measure over the receiver operator characteristic curves (Fawcett, 2006) of each model when applied to the leave-out sample. The AUC measure takes a value of 0.5 for noninformative predictors, and values above it and up to 1 indicate predicting power in binary classification problems. Both models have AUC close to 0.6, but the reputation model has a significantly larger AUC than the popularity model, as evidenced in a DeLong test ($Z = 20.9576$, $p < 10^{-15}$). This evidence supports the additional value of global

reputation metrics in predicting, when users will become inactive, as stated in the cascading inactivity hypothesis.

The above findings are robust to a series of diagnostic and robustness tests, as presented in Appendix SII. To cope with the correlations in high positive values introduced by polynomial terms, we repeated the estimate using orthogonalized polynomials (Kennedy & Gentle, 1980). The shape of the main predictors and predictive power were unaffected, showing that the results of the fit are robust to correlations introduced by the polynomials. Additionally, we tested if collinearity between popularity and reputation could change the shape of the predictions, by analyzing the residuals of partial models over the leave-out sample. We found the same nonmonotonic shapes, when analyzing the dependence of the residualized probability of inactivity versus reputation and popularity, illustrating that our results are robust to correlations between reputation and popularity.

Popularity Analysis

To understand changes in popularity, we focus on the set of users that were active since April 1, 2016, and exclude a random 10 percent of those users as a leave-out sample. Figure 4 shows a histogram of the amounts of followers in 2009 and 2016. As expected given the growth of Twitter as a social network, popularity in terms of amount of followers had a tendency to grow between the two time points. This growth is especially present for users of low to moderate popularity, as the major density of the histogram is significantly above the diagonal, but close to the axes. When going to larger values of popularity, growth seems to be moderated, suggesting a sublinear yet powerful growth tendency.

We inspect the distribution of popularity as measured by the amount of followers, to formulate a model to predict popularity in 2016 based on the full 2009 network. The left panel of Figure 5 shows the density function of the

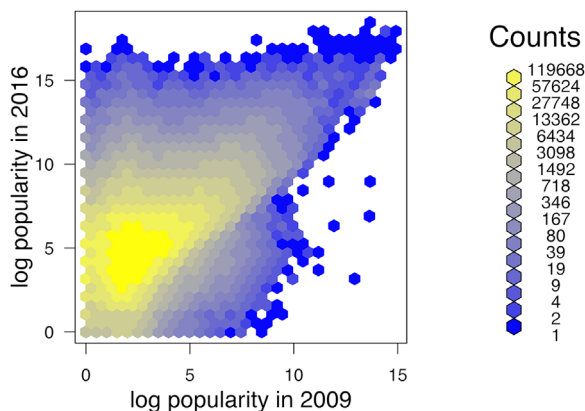


Figure 4. Histogram of Logarithms of the Amount of Followers as Measured in 2009 Versus 2016, With Frequency in the Bins on a Logarithmic Scale.

Note: The vast majority of the users have a large significant increase in followers, with a tendency to be saturated over very popular users.




Figure 5. Left: Probability density function of the logarithm of the amount of followers of active users in 2016. Right: Complementary cumulative density function (CCDF) of popularity of active users with power-law fit (blue) and log-normal fit (red). The distribution is best fit by a log-normal, outperforming a power-law.

logarithm of the amount of followers. The amount of followers closely follows a log-normal distribution, with a distribution of logarithms close to normality. Some skewness can be seen on the density function, which can be attributed to the error associated to discrete values.

We evaluate the evidence for the log-normal hypothesis versus the alternative and more common power-law hypothesis. We apply maximum likelihood estimators based on a tail-reweighted Kolmogorov–Smirnov criterion (Clauset et al., 2009), identifying the minimum value of followers (368) over which a power-law tail might be present. We then performed the same tail estimate using a log-normal distribution instead, and the result of both fits is shown against the complementary cumulative density function on the right panel of Figure 5. The fit of a log-normal distribution outperforms a power-law fit, reaching a lower Kolmogorov–Smirnov distance below 0.1. A two-sided likelihood ratio test (Vuong, 1989) supports this result: The log-likelihood ratio between the power-law model and the log-normal one is -17.1 ($p < 10^{-60}$). From this, we can conclude that, while the tail of popularity is rather long, the popularity of early adopters that are still active is more likely to follow a log-normal distribution than a power-law.

Given the log-normality of the amount of followers, we model the patterns observed over the 2D histogram as a linear regression to predict the logarithm of the amount of followers in 2016, based on the full 2009 network metrics. We focus on users with at least one follower in 2009 and 2016, including a total of more than 5.4 million users. First, we model how the growth in the amount of followers is mediated by reputation:

$$F_{16} = a_I + b_I \cdot F_{09} + c_I \cdot I_{09} + d_I \cdot F_{09} \cdot I_{09} + \epsilon_I \quad (2)$$

where F_{16} and F_{09} are the logarithm of the amount of followers in 2016 and 2009, respectively, and I_{09} is the logarithm of the in-coreness in 2009. The error term ϵ_I is assumed to be uncorrelated and normally distributed.

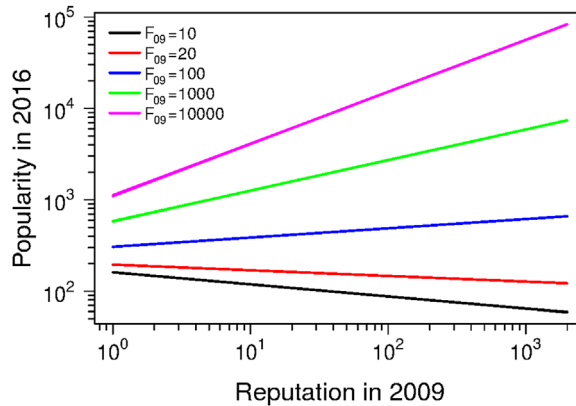


Figure 6. Model Predictions for Popularity in 2016 Versus Reputation, After Fixing the Value of Popularity in 2009 to a Range of Values.

Notes: Reputation has a positive effect on the future popularity of a user when users are popular enough, otherwise it decreases it. Shaded areas shows standard errors and are too small to be visible.

We find that popularity in 2009 leads to higher popularity in 2016, as evidenced by the positive estimate of b_F , as shown in Appendix SIII along with more detailed fit results. The interaction with reputation is not trivial, the base effect of reputation is negative, but the interaction term with popularity is positive. To understand it better, we plot the predicted value of popularity in 2016 versus reputation in 2009, for a selection of various values of popularity in 2009.

As shown on the Figure 6, for low values of popularity in 2009, the effect of reputation is slightly negative. After popularity values above 100 followers, the effect of reputation is positive, showing that reputation helps to growth popularity of already popular users, while it is not a necessarily good resource for less popular ones. These results lend support for the popularity growth hypothesis, and show an interesting mixed scenario for the weak competition hypothesis: reputation can have a positive effect on popularity, but cannot replace it.

To understand the effect of the bow-tie structure on the changes in popularity, we follow a similar modeling approach as with reputation:

$$F_{16} = a_B + b_B \cdot F_{09} + c_B \cdot B_{09} + d_B \cdot F_{09} \cdot B_{09} + \epsilon_B \quad (3)$$

where B_{09} is a categorical factor that captures if the user was in the In, Out, or SCC group in 2009. The results of the fit reveal significant effects of the bow-tie group in future popularity, presented more in detail in Appendix SIII. Figure 7 shows the predicted value of popularity in 2016 versus 2009, for the four bow-tie groups in the fitted model. Growth in popularity is generally sublinear, and the tendency to grow depends on the bow-tie group. Users that belonged to the SCC and the Out group grew the most in popularity in comparison to other groups.

We evaluate the generalizability of these results predicting popularity in the leave-out sample. A null model with only an intercept has a mean absolute error (MAE) of

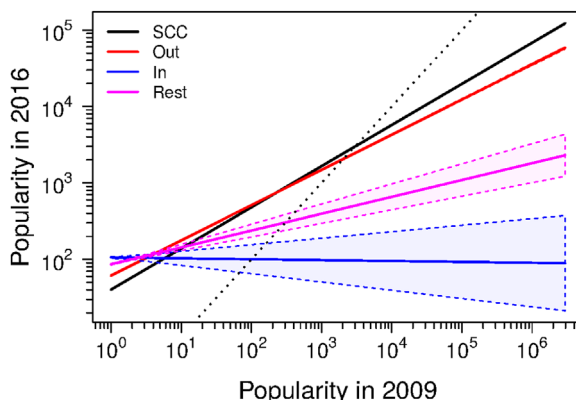


Figure 7. Model Predictions of the Growth of Popularity for the Four Bow-Tie Groups, With the Identity Dashed Line as Reference.

Notes: Growth is evident but sublinear in the SCC and Out groups, moderated in the rest, and not significant in the In group. Shaded areas show standard errors around the prediction.

1.370. A model using only popularity in 2009 has a MAE of 1.242, which is improved when adding reputation to 1.220 and when adding the bow-tie group to 1.239.

Social Influence Analysis

After testing the role of reputation and popularity in the tendency of users to become inactive and grow in popularity, we explore how reputation and popularity shape social influence. As with the other analyses, we remove 10 percent of users as a leave-out sample, including a total of 3,300,045 users in the fits.

The division of impact hypothesis postulates that social influence will grow as a power function of the size of the group, that is, the amount of followers, where the exponent is below 1. To test this hypothesis, we compute S_{RT} as the logarithm of the average amount of retweets, and formulate a regression model as:

$$S_{RT} = a_{RT} + b_{RT} \cdot F_{09} + c_{RT} \cdot I_{09} + d_{RT} \cdot F_{09} \cdot I_{09} + \epsilon_{RT} \tag{4}$$

where F_{09} and I_{09} are the logmodulus of popularity and reputation in 2009, respectively. Equation (4) models the logarithm of social influence as a linear combination of F_{09} , I_{09} , and their interaction.

Figure 8 shows the predicted value of the average amount of retweets versus popularity for the fit, as explained more in detail in Appendix SIV along with coefficient estimates and regression diagnostics. It can be seen that, along various values of reputation, social influence scales sublinearly with popularity. More precisely, the fit estimates $b_{RT}=0.54$ with a standard error below 0.01, an estimate clearly below 1 as hypothesized. Due to the positive interaction term, this relationship approaches linearity along the reputation scale, as appreciated in the line for extremely high reputation in Figure 8. This is exceptional in the range of




Figure 8. Predicted Value of Social Influence in the Fit of the Model of Equation (4).
 Notes: Social influence is shown versus popularity in 2009 for various values of reputation (I_{09}). Social influence grows sublinearly with popularity and the growth approaches linearity for higher reputation.

values observed in our analysis, suggesting more heterogeneous samples might be able to find positive effects of reputation.

To better understand how popularity and reputation interact in social influence, we show the prediction profile for the model in the left panel Figure 9. Social influence grows with popularity but not with reputation, which shows a soft decreasing effect. This negative effects of reputation softens as popularity increases, but never reaches a positive trend in the observed value range.

A null model has an MAE over the leave out sample of 1.243, outperformed by the MAE of 1.226 of the model of Equation (4) (t -test $p < 10^{-5}$). This model outperforms models that use only popularity ($p < 10^{-5}$) and only reputation ($p < 10^{-5}$). Comparing the leave-out sample errors of models using only popularity and only reputation reveals that the popularity model is a better fit ($p < 10^{-5}$), an observation that is also supported by a likelihood ratio test over the

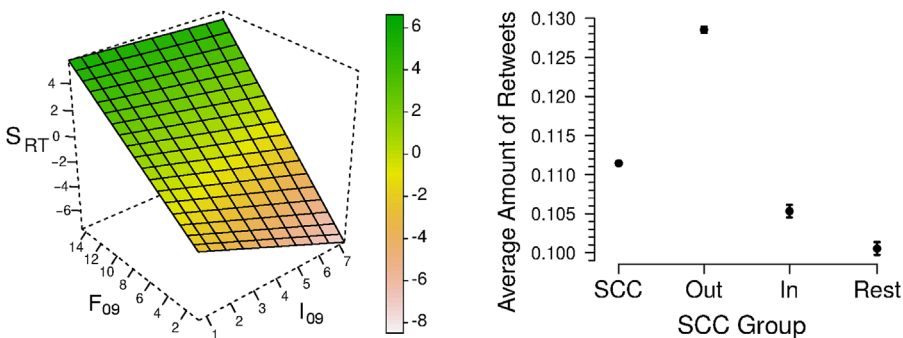


Figure 9. Left: Predicted value of social influence over the range of popularity and reputation in 2009. Social influence grows with popularity, and decreases with reputation when popularity is low. Right: Predicted values of social influence for users in the bow-tie groups of the network. Users in the Out group are significantly more influential than the rest.

training data. This evidence runs against the reputation strength hypothesis, suggesting that popularity is a better predictor for social influence in Twitter.

We test the bow-tie influence hypothesis with a model similar to the one of Equation (4), but with a dependent categorical variable with the bow-tie group as in Equation (3). The detailed results of the fit and its diagnostics are reported in Appendix SIV, and the right panel of Figure 9 shows the predicted values of the average amount of retweets for the four groups of the bow-tie decomposition. This analysis reveals that users in the Out group are significantly more influential than other users, supporting the bow-tie influence hypothesis. In relative terms, a user in the Out group is 15.3 percent more influential than one in the SCC, and 27.9 percent more influential than one in the Rest group.

Discussion

We present a large-scale analysis of the relations between popularity, reputation, activity, and social influence across a period of seven years in Twitter. We processed the full 2009 network of followers, calculating popularity and reputation values for more than 40 million Twitter users, and revealing a bow-tie structure with a giant strongly connected component. Our results can be summarized as follows.

1. Results on activity

- i *Nonmonotonic motivation hypothesis: Supported.* Logistic regression results show that the probability of inactivity has a generally negative trend, but nonlinear polynomial terms are significant and show local regions of positive trends.
- ii *Cascading inactivity hypothesis: Supported.* The MAE in the leave-out sample for a logistic regression model using only reputation is lower than using only popularity.
- iii *Core activity hypothesis: Supported.* Tests of equality of proportions of becoming inactive in the bow-tie groups show that the probability of becoming inactive in the SCC is lower than in the other four parts.

2. Results on popularity

- i *Power-law popularity hypothesis: Rejected.* A log-normal distribution fit of popularity outperforms a power-law, even when fitting only the tail of the distribution.
- ii *Popularity growth hypothesis: Supported.* A regression model shows that popularity in 2016 is sublinearly and positively associated with popularity in 2009, even when controlling for reputation.
- iii *Weak competition hypothesis: Partially supported.* Regression results show a negative association between reputation in 2009 and popularity in 2016, but

the interaction term is positive. As a result, reputation has a positive effect on future popularity only for sufficiently high levels of popularity in 2009.

3. Results on social influence

- i *Division of impact hypothesis: Supported.* A linear model fit shows that average retweet rates grow sublinearly with popularity in 2009.
- ii *Bow-tie influence hypothesis: Supported.* The retweet rate of users in the Out group of the bow-tie decomposition are significantly higher on average.
- iii *Reputation strength: Rejected.* The MAE in the leave-out sample for a regression model of retweet rates using only reputation is higher than for a model using only popularity.

Our analysis of activity in 2016 shows that the probability to become inactive in Twitter follows a nonmonotonic function with respect to popularity and reputation. This contrasts with the initial assumption of resilience models for online social networks (Garcia et al., 2013): a monotonically decreasing inactivity probability with degree. Our results show that a general negative trend is in place, but also show a concavity pattern that calls for new microdynamic formulations, as suggested in the framework of Garcia et al. (2013) and recently explored by Yu et al. (2016). Nevertheless, we must note that the cascading inactivity hypothesis follows from these kind of models, which we find support for when comparing the quality of predictors based on degree and coreness. This is an example of how simple microscopic assumptions, which might be incomplete or even empirically incorrect, can be used to formulate hypotheses about macroscopic phenomena, when analyzed as an agent-based model (Ball, 2007).

We find that users outside the SCC have a significant increased probability of becoming inactive, thus being excluded from social and political phenomena present in Twitter. This can be seen an example of the second digital divide (Attewell, 2001), where the difference in participation is not due to inequality in the access to information technologies, but on the variance of intensity and proficiency of their use. As such, the citizens that could be excluded from deliberation in Twitter are not the ones that could not have access to a computer or smartphone, but the ones that could not afford the time or effort to keep active and engaged in the Twitter society. An alternative interpretation to this finding lies on the large size of the SCC: the larger activity level might be a size-induced phenomenon rather than being generated by the connections within the SCC. Our analysis cannot address the causation mechanism for this effect of the SCC on activity, in particular whether it is an effect on motivation or a selection mechanism that keeps fake accounts or spammers outside the SCC. Only future comparative analyses with other online social networks and with more detailed data will allow us to understand why the bow-tie structure is linked to user inactivity.

Our analysis of the distribution of popularity clearly refutes the power-law hypothesis, supporting a log-normal distribution as a better explanation. A theoretical explanation for the log-normal hypothesis can be found on stochastic growth

models. Power-law distributions emerge from a combination of a birth process and proportional growth (Mitzenmacher, 2004), while log-normal distributions are a conclusion of proportional growth alone. In our study, this can emerge from the difference between the time scale of users joining Twitter, roughly three years for this sample, when compared to the time scale of their growth in popularity, which is seven years. The difference between these time scales implies that the birth process of users is negligible in comparison to the proportional growth of the amounts of followers, and thus the final distribution resembles one in which all users would have joined Twitter roughly on the same date: a log-normal distribution.

When predicting popularity as the amount of followers in 2016, we find that popularity grows with reputation only for users with some previous level of popularity. Thus, with respect to the growth in attention as a resource, reputation is a complementary but costly resource that seems to pay off only in the case of sufficiently popular users. Nevertheless, this points to a weak competition scenario in which reputation is not detrimental and being followed by reputable users does not have a negative effect on future popularity.

Our analysis of social influence shows that the extent of social influence grows sublinearly with popularity, supporting the division of impact hypothesis drawn from social impact theory. Interestingly, we find that popularity alone is a better predictor of social influence than reputation, refuting the reputation strength hypothesis. On the one hand, this contrasts with previous research that compared the predictive power of undirected coreness metrics and degree, finding in pairwise comparisons that coreness is a better predictor of spreading in simulations (Kitsak et al., 2010) and empirical data (Pei, Muchnik, Andrade, Zheng, & Makse, 2014). On the other hand, information spreading in social network has been repeatedly shown to greatly differ from epidemic spreading (Lerman, 2016). This calls for a future study that can disentangle, whether the discrepancy between simulation results and our empirical analysis is due to topological properties of the follower network or due to complex dynamics in retweeting behavior.

Our work suffers the current research bias of Twitter as a model organism (Tufekci, 2014): The generalization of results is limited due to the large amount of research only based on Twitter data. This overfocusing on Twitter carries some risks for the scientific community as a whole, for example in relation with the Terms of Service (Watters, 2011). We chose Twitter for our analysis because of its large size and longitudinally tractable nature, two properties that are not always available for public research in other social networks. Our perspective measuring popularity through the amount of followers and reputation through in-coreness is extensible to other directed social networking sites such as Instagram or Youtube, and thus future research can explore if other sites follow similar mechanics as the ones we revealed here. Furthermore, fine grained takes on user behavior can study important aspects of popularity, for example, on different classes of accounts related to companies and promotional content (Abisheva et al., 2014), or on linguistic and sentiment signals in popularity growth (Niven et al., 2015).

Our research has followed a theory-driven approach, aiming at aspects of Twitter that are relevant for society as a whole and that have been inferred from

both qualitative and computational models. This allowed us take a step further from descriptive and data-driven analyses without theoretical context, framing our findings in a wider scientific perspective beyond the computational sciences (Cihon & Yasserli, 2016). We learned that popularity and reputation are not always motivating, that popularity is not as heterogeneous as was thought to be, and that popularity and reputation are both relevant when studying social influence. Our methods can be applied to subsets of Twitter in a country, detecting influential individuals, finding excluded social groups that become inactive, and predicting the heterogeneous potential that social media has when spreading information.

David Garcia, Ph.D., Chair of Systems Design, ETH Zurich, Weinbergstrasse 56/58, 8092 Zurich, Switzerland [dgarcia@ethz.ch].

Pavlin Mavrodiev, Ph.D., Chair of Systems Design, ETH Zurich, Weinbergstrasse 56/58, 8092 Zurich, Switzerland.

Daniele Casati, M.Sc., Chair of Systems Design, ETH Zurich, Weinbergstrasse 56/58, 8092 Zurich, Switzerland.

Frank Schweitzer, Prof., Chair of Systems Design, ETH Zurich, Weinbergstrasse 56/58, 8092 Zurich, Switzerland.

Notes

This was funded by the ETH Risk Centre's "Systemic Risks for Privacy in Online Interaction" and by Schweizerischer Nationalfonds zur Förderung der Wissenschaftlichen Forschung (Grant number: CR2111_146499).

1. <https://an.kaist.ac.kr/traces/WWW2010.html>.
2. We performed a second request on all missing users and found no result on any of them. This discards programmatic errors as a possible source for the missing users.

References

- Abisheva, A., V.R.K. Garimella, D. Garcia, and I. Weber. 2014. "Who Watches (and Shares) What on Youtube? And When? Using Twitter to Understand Youtube Viewership." In *Proceedings of the 7th ACM International Conference on Web Search and Data Mining*. New York, NY: ACM, 593–602.
- Alstott, J., E. Bullmore, and D. Plenz. 2014. "Powerlaw: A Python Package for Analysis of Heavytailed Distributions." *PLoS One* 9 (1): e85777.
- Alvarez, R., D. Garcia, Y. Moreno, and F. Schweitzer. 2015. "Sentiment Cascades in the 15M Movement." *EPJ Data Science* 4 (1): 1–13.
- Aragón, P., K.E. Kappler, A. Kaltenbrunner, D. Laniado, and Y. Volkovich. 2013. "Communication Dynamics in Twitter During Political Campaigns: The Case of the 2011 Spanish National Election." *Policy & Internet* 5 (2): 183–206.
- Attewell, P. 2001. "Comment: The First and Second Digital Divides." *Sociology of Education* 74 (3): 252–59.
- Ball, P. 2007. "Social Science Goes Virtual." *Nature* 448 (7154): 647–48.
- Barberá, P. 2015. "Birds of the Same Feather Tweet Together: Bayesian Ideal Point Estimation Using Twitter Data." *Political Analysis* 23 (1): 76–91.
- Bennett, W.L., and A. Segerberg. 2012. "The Logic of Connective Action: Digital Media and the Personalization of Contentious Politics." *Information, Communication, & Society* 15 (5): 739–68.
- Borge-Holthoefer, J., and Y. Moreno. 2012. "Absence of Influential Spreaders in Rumor Dynamics." *Physical Review E* 85(2): 026116.

- boyd, d., and K. Crawford. 2012. "Critical Questions for Big Data: Provocations for a Cultural, Technological, and Scholarly Phenomenon." *Information, Communication, & Society* 15 (5): 662–79.
- Broder, A., R. Kumar, F. Maghoul, P. Raghavan, S. Rajagopalan, R. Stata, A. Tomkins, and J. Wiener. 2000. "Graph Structure in the Web." *Computer Networks* 33 (1): 309–20.
- Castells, M. 2011. *The Rise of the Network Society: The Information Age: Economy, Society, and Culture*, Vol. 1. Oxford, UK: John Wiley & Sons.
- Cihon, P., and T. Yasseri. 2016. "A Biased Review of Biases in Twitter Studies on Political Collective Action." *Frontiers in Physics* 4: 34.
- Clauset, A., C.R. Shalizi, and M.E. Newman. 2009. "Power-Law Distributions in Empirical Data." *SIAM Review* 51 (4): 661–703.
- Conover, M.D., B. Gonçalves, A. Flammini, and F. Menczer. 2012. "Partisan Asymmetries in Online Political Activity." *EPJ Data Science* 1 (1): 1.
- Fawcett, T. 2006. "An Introduction to ROC Analysis." *Pattern Recognition Letters* 27 (8): 861–74.
- Freeman, L.C. 1977. "A Set of Measures of Centrality Based on Betweenness." *Sociometry* 40 (1): 35–41.
- Friedkin, N.E. 1991. "Theoretical Foundations for Centrality Measures." *American Journal of Sociology* 96 (6): 1478–504.
- Garcia, D., P. Mavrodiev, and F. Schweitzer. 2013. "Social Resilience in Online Communities: The Autopsy of Friendster." In *Proceedings of the First ACM Conference on Online Social Networks*. New York, NY: ACM, 39–50.
- Gelman, A., A. Jakulin, M.G. Pittau, and Y.-S. Su. 2008. "A Weakly Informative Default Prior Distribution for Logistic and Other Regression Models." *The Annals of Applied Statistics* 2 (4): 1360–83.
- Giatsidis, C., D.M. Thilikos, and M. Vazirgiannis. 2013. "D-Cores: Measuring Collaboration of Directed Graphs Based on Degeneracy." *Knowledge and Information Systems* 35 (2): 311–43.
- González-Bailón, S. 2013. "Social Science in the Era of Big Data." *Policy & Internet* 5 (2): 147–60.
- González-Bailón, S., J. Borge-Holthoefer, A. Rivero, and Y. Moreno. 2011. "The Dynamics of Protest Recruitment Through an Online Network." *Scientific Reports* 1. Article number: 197. <http://www.nature.com/articles/srep00197>.
- Guille, A., H. Hacid, C. Favre, and D.A. Zighed. 2013. "Information Diffusion in Online Social Networks: A Survey." *ACM SIGMOD Record* 42 (2): 17–28.
- Hiltz, S.R., and M. Turoff. 1985. "Structuring Computer-Mediated Communication Systems to Avoid Information Overload." *Communications of the ACM* 28 (7): 680–89.
- Huberman, B.A., and F. Wu. 2008. "The Economics of Attention: Maximizing User Value in Information-Rich Environments." *Advances in Complex Systems* 11 (4): 487–96.
- Jain, S., and S. Krishna. 2001. "A Model for the Emergence of Cooperation, Interdependence, and Structure in Evolving Networks." *Proceedings of the National Academy of Sciences* 98 (2): 543–47.
- Kennedy, W.J.W.J., and J.E. Gentle. 1980. *Statistical Computing*. <https://www.amazon.com/Statistical-Computing-W-Kennedy/dp/0824784065>.
- Kitsak, M., L.K. Gallos, S. Havlin, F. Liljeros, L. Muchnik, H.E. Stanley, and H.A. Makse. 2010. "Identification of Influential Spreaders in Complex Networks." *Nature Physics* 6 (11): 888–93.
- Kleinberg, J.M. 1999. "Hubs, Authorities, and Communities." *ACM Computing Surveys (CSUR)* 31(4es): 5.
- Kwak, H., C. Lee, H. Park, and S. Moon. 2010. "What is Twitter, a Social Network or a News Media?" In *Proceedings of the 19th International Conference on World Wide Web*. New York, NY: ACM, 591–600.
- Latané, B. 1981. "The Psychology of Social Impact." *American Psychologist* 36 (4): 343–56.
- Lazer, D., A.S. Pentland, L. Adamic, S. Aral, A.L. Barabasi, D. Brewer, N. Christakis, et al. 2009. "Life in the Network: The Coming Age of Computational Social Science." *Science* 323 (5915): 721.
- Lerman, K. 2016. "Information Is Not a Virus, and Other Consequences of Human Cognitive Limits." *Future Internet* 8(2): 21.
- Lietz, H., C. Wagner, A. Bleier, and M. Strohmaier. 2014. "When Politicians Talk: Assessing Online Conversational Practices of Political Parties on Twitter." In *International AAAI Conference on Web and Social Media*. Ann Arbor, MI: AAAI.

- MacPhee, W.N. 1963. *Formal Theories of Mass Behaviour*. New York: Free Press of Glencoe, 244.
- Mitzenmacher, M. 2004. "A Brief History of Generative Models for Power Law and Lognormal Distributions." *Internet Mathematics* 1 (2): 226–51.
- Niven, K., D. Garcia, I. van der Löwe, D. Holman, and W. Mansell. 2015. "Becoming Popular: Interpersonal Emotion Regulation Predicts Relationship Formation in Real Life Social Networks." *Frontiers in Psychology* 6: 1452.
- Olson, M. 1965. *The Logic of Collective Action*, Vol. 124. Cambridge, MA: Harvard University Press.
- Pei, S., L. Muchnik, J.S. Andrade Jr., Z. Zheng, and H.A. Makse. 2014. "Searching for Superspreaders of Information in Real-World Social Media." *Scientific Reports* 4. Article number: 5547. <http://www.nature.com/articles/srep05547>.
- Riquelme, F., and P. González-Cantergiani. 2016. "Measuring User Influence on Twitter: A Survey." *Information Processing & Management* 52 (5): 949–75.
- Rodriguez, M.G., K. Gummadi, and B. Schoelkopf. 2014. "Quantifying Information Overload in Social Media and Its Impact on Social Contagions." In *International AAAI Conference on Web and Social Media*. Ann Arbor, MI: AAAI.
- Seidman, S.B. 1983. "Network Structure and Minimum Degree." *Social Networks* 5 (3): 269–87.
- Stumpf, M.P., and M.A. Porter. 2012. "Critical Truths About Power Laws." *Science* 335 (6069): 665–66.
- Toubia, O., and A.T. Stephen. 2013. "Intrinsic vs. Image-Related Utility in Social Media: Why Do People Contribute Content to Twitter?" *Marketing Science* 32 (3): 368–92.
- Tufekci, Z., 2014. "Big Questions for Social Media Big Data: Representativeness, Validity and Other Methodological Pitfalls." *ICWSM*. <http://www.aaai.org/ocs/index.php/ICWSM/ICWSM14/paper/view/8062>.
- Vitali, S., J.B. Glattfelder, and S. Battiston. 2011. "The Network of Global Corporate Control." *PLoS One* 6 (10): e25995.
- Vuong, Q.H. 1989. "Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses." *Econometrica: Journal of the Econometric Society* 57 (2): 307–33.
- Watters, A. 2011. "How Recent Changes to Twitter's Terms of Service Might Hurt Academic Research." *Read Write*. http://readwrite.com/2011/03/03/how_recent_changes_to_twiters_terms_of_service_mi/.
- Webb, E.J., D.T. Campbell, R.D. Schwartz, and L. Sechrest. 1966. *Unobtrusive Measures: Nonreactive Research in the Social Sciences*, Vol. 111. Chicago, IL: Rand McNally Chicago.
- Wilson, E.B. 1927. "Probable Inference, the Law of Succession, and Statistical Inference." *Journal of the American Statistical Association* 22 (158): 209–12.
- Yu, Y., G. Xiao, J. Zhou, Y. Wang, Z. Wang, J. Kurths, and H.J. Schellnhuber. 2016. "System Crash as Dynamics of Complex Networks." *Proceedings of the National Academy of Sciences* 113 (42): 11726–31.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's web-site.

Appendix SI. Bow-tie group and inactivity.

Appendix SII. Inactivity model fits.

Appendix SIII. Popularity model fits.

Appendix SIV. Social influence model fits.