The spreading of misinformation online

Michela Del Vicario[a], Alessandro Bessi[b], Fabiana Zollo[a], Fabio Petroni[c], Antonio Scala[a,d], Guido Caldarelli[a,d], H. Eugene Stanley[e], and Walter Quattrociocchi[a,1]

[a]Laboratory of Computational Social Science, Networks Department, IMT Alti Studi Lucca, 55100 Lucca, Italy; [b]IUSS Institute for Advanced Study, 27100 Pavia, Italy; [c]Sapienza University, 00185 Rome, Italy; [d]ISC-CNR Uos "Sapienza," 00185 Rome, Italy; and [e]Boston University, Boston, MA 02115

Edited by Matjaz Perc, University of Maribor, Maribor, Slovenia, and accepted by the Editorial Board December 4, 2015 (received for review September 1, 2015)

The wide availability of user-provided content in online social media facilitates the aggregation of people around common interests, worldviews, and narratives. However, the World Wide Web (WWW) also allows for the rapid dissemination of unsubstantiated rumors and conspiracy theories that often elicit rapid, large, but naive social responses such as the recent case of Jade Helm 15—where a simple military exercise turned out to be perceived as the beginning of a new civil war in the United States. In this work, we address the determinants governing misinformation spreading through a thorough quantitative analysis. In particular, we focus on how Facebook users consume information related to two distinct narratives: scientific and conspiracy news. We find that, although consumers of scientific and conspiracy stories present similar consumption patterns with respect to content, cascade dynamics differ. Selective exposure to content is the primary driver of content diffusion and generates the formation of homogeneous clusters, i.e., "echo chambers." Indeed, homogeneity appears to be the primary driver for the diffusion of contents and each echo chamber has its own cascade dynamics. Finally, we introduce a data-driven percolation model mimicking rumor spreading and we show that homogeneity and polarization are the main determinants for predicting cascades' size.

misinformation | virality | Facebook | rumor spreading | cascades

The massive diffusion of sociotechnical systems and microblogging platforms on the World Wide Web (WWW) creates a direct path from producers to consumers of content, i.e., allows disintermediation, and changes the way users become informed, debate, and form their opinions (1–5). This disintermediated environment can foster confusion about causation, and thus encourage speculation, rumors, and mistrust (6). In 2011 a blogger claimed that global warming was a fraud designed to diminish liberty and weaken democracy (7). Misinformation about the Ebola epidemic has caused confusion among healthcare workers (8). Jade Helm 15, a simple military exercise, was perceived on the Internet as the beginning of a new civil war in the United States (9).

Recent works (10–12) have shown that increasing the exposure of users to unsubstantiated rumors increases their tendency to be credulous.

According to ref. 13, beliefs formation and revision is influenced by the way communities attempt to make sense of events or facts. Such a phenomenon is particularly evident on the WWW where users, embedded in homogeneous clusters (14–16), process information through a shared system of meaning (10, 11, 17, 18) and trigger collective framing of narratives that are often biased toward self-confirmation.

In this work, through a thorough quantitative analysis on a massive dataset, we study the determinants behind misinformation diffusion. In particular, we analyze the cascade dynamics of Facebook users when the content is related to very distinct narratives: conspiracy theories and scientific information. On the one hand, conspiracy theories simplify causation, reduce the complexity of reality, and are formulated in a way that is able to tolerate a certain level of uncertainty (19–21). On the other hand, scientific information disseminates scientific advances and exhibits the process of scientific thinking. Notice that we do not focus on the quality of the information but rather on the possibility of verification. Indeed, the main difference between the two is content verifiability. The generators of scientific information and their data, methods, and outcomes are readily identifiable and available. The origins of conspiracy theories are often unknown and their content is strongly disengaged from mainstream society and sharply divergent from recommended practices (22), e.g., the belief that vaccines cause autism.

Massive digital misinformation is becoming pervasive in online social media to the extent that it has been listed by the World Economic Forum (WEF) as one of the main threats to our society (23). To counteract this trend, algorithmic-driven solutions have been proposed (24–29), e.g., Google (30) is developing a trustworthiness score to rank the results of queries. Similarly, Facebook has proposed a community-driven approach where users can flag false content to correct the newsfeed algorithm. This issue is controversial, however, because it raises fears that the free circulation of content may be threatened and that the proposed algorithms may not be accurate or effective (10, 11, 31). Often conspiracists will denounce attempts to debunk false information as acts of misinformation.

Whether a claim (either substantiated or not) is accepted by an individual is strongly influenced by social norms and by the claim's coherence with the individual's belief system—i.e., confirmation bias (32, 33). Many mechanisms animate the flow of false information that generates false beliefs in an individual, which, once adopted, are rarely corrected (34–37).

In this work we provide important insights toward the understanding of cascade dynamics in online social media and in particular about misinformation spreading.

We show that content-selective exposure is the primary driver of content diffusion and generates the formation of homogeneous

Significance

The wide availability of user-provided content in online social media facilitates the aggregation of people around common interests, worldviews, and narratives. However, the World Wide Web is a fruitful environment for the massive diffusion of unverified rumors. In this work, using a massive quantitative analysis of Facebook, we show that information related to distinct narratives—conspiracy theories and scientific news—generates homogeneous and polarized communities (i.e., echo chambers) having similar information consumption patterns. Then, we derive a data-driven percolation model of rumor spreading that demonstrates that homogeneity and polarization are the main determinants for predicting cascades' size.

Author contributions: M.D.V., A.B., F.Z., A.S., G.C., H.E.S., and W.Q. designed research; M.D.V., A.B., F.Z., H.E.S., and W.Q. performed research; M.D.V., A.B., F.Z., F.P., and W.Q. contributed new reagents/analytic tools; M.D.V., A.B., F.Z., A.S., G.C., H.E.S., and W.Q. analyzed data; and M.D.V., A.B., F.Z., A.S., G.C., H.E.S., and W.Q. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission. M.P. is a guest editor invited by the Editorial Board.

Freely available online through the PNAS open access option.

[1]To whom correspondence should be addressed. Email: walterquattrociocchi@gmail.com.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1517441113/-/DCSupplemental.

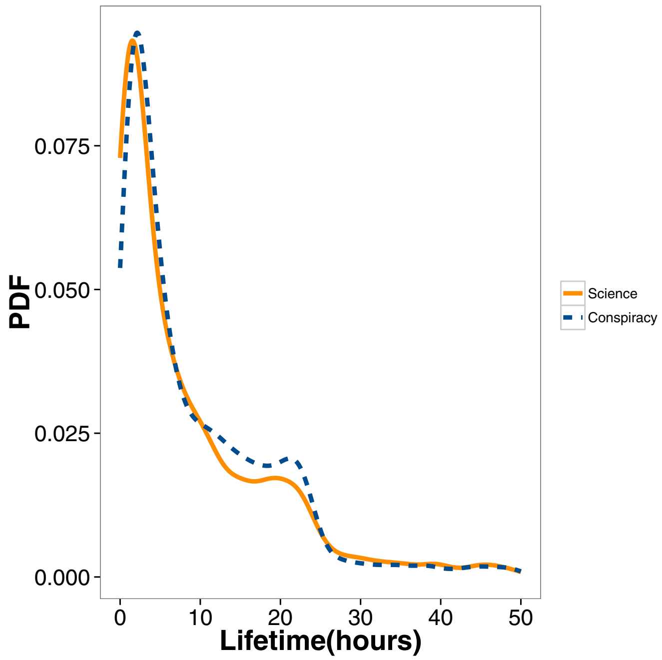

Fig. 1. PDF of lifetime computed on science news and conspiracy theories, where the lifetime is here computed as the temporal distance (in hours) between the first and last share of a post. Both categories show a similar behavior.

clusters, i.e., "echo chambers" (10, 11, 38, 39). Indeed, our analysis reveals that two well-formed and highly segregated communities exist around conspiracy and scientific topics. We also find that although consumers of scientific information and conspiracy theories exhibit similar consumption patterns with respect to content, the cascade patterns of the two differ. Homogeneity appears to be the preferential driver for the diffusion of content, yet each echo chamber has its own cascade dynamics. To account for these features we provide an accurate data-driven percolation model of rumor spreading showing that homogeneity and polarization are the main determinants for predicting cascade size.

The paper is structured as follows. First we provide the preliminary definitions and details concerning data collection. We then provide a comparative analysis and characterize the statistical signatures of cascades of the different kinds of content. Finally, we introduce a data-driven model that replicates the analyzed cascade dynamics.

Methods

Ethics Statement. Approval and informed consent were not needed because the data collection process has been carried out using the Facebook Graph application program interface (API) (40), which is publicly available. For the analysis (according to the specification settings of the API) we only used publicly available data (thus users with privacy restrictions are not included in the dataset). The pages from which we download data are public Facebook entities and can be accessed by anyone. User content contributing to these pages is also public unless the user's privacy settings specify otherwise, and in that case it is not available to us.

Data Collection. Debate about social issues continues to expand across the Web, and unprecedented social phenomena such as the massive recruitment of people around common interests, ideas, and political visions are emerging. Using the approach described in ref. 10, we define the space of our investigation with the support of diverse Facebook groups that are active in the debunking of misinformation.

The resulting dataset is composed of 67 public pages divided between 32 about conspiracy theories and 35 about science news. A second set, composed of two troll pages, is used as a benchmark to fit our data-driven model. The first category (conspiracy theories) includes the pages that disseminate alternative, controversial information, often lacking supporting evidence and frequently advancing conspiracy theories. The second category (science news) includes the pages that disseminate scientific information. The third category (trolls) includes those pages that intentionally disseminate sarcastic false information on the Web with the aim of mocking the collective credulity online.

For the three sets of pages we download all of the posts (and their respective user interactions) across a 5-y time span (2010–2014). We perform the data collection process by using the Facebook Graph API (40), which is publicly available and accessible through any personal Facebook user account. The exact breakdown of the data is presented in *SI Appendix*, section 1.

Preliminaries and Definitions. A tree is an undirected simple graph that is connected and has no simple cycles. An oriented tree is a directed acyclic graph whose underlying undirected graph is a tree. A sharing tree, in the context of our research, is an oriented tree made up of the successive sharing of a news item through the Facebook system. The root of the sharing tree is the node that performs the first share. We define the size of the sharing tree as the number of nodes (and hence the number of news sharers) in the tree and the height of the sharing tree as the maximum path length from the root.

We define the user polarization $\sigma = 2\varrho - 1$, where $0 \le \varrho \le 1$ is the fraction of "likes" a user puts on conspiracy-related content, and hence $-1 \le \sigma \le 1$. From user polarization, we define the edge homogeneity, for any edge e_{ij} between nodes i and j, as

$$\sigma_{ij} = \sigma_i \sigma_j,$$

with $-1 \le \sigma_{ij} \le 1$. Edge homogeneity reflects the similarity level between the polarization of the two sharing nodes. A link in the sharing tree is

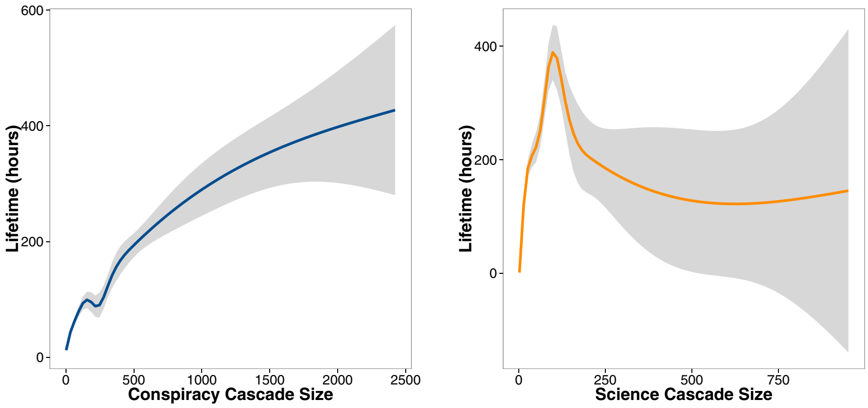

Fig. 2. Lifetime as a function of the cascade size for conspiracy news (*Left*) and science news (*Right*). Science news quickly reaches a higher diffusion; a longer lifetime does not correspond to a higher level of interest. Conspiracy rumors are assimilated more slowly and show a positive relation between lifetime and size.

homogeneous if its edge homogeneity is positive. We then define a sharing path to be any path from the root to one of the leaves of the sharing tree. A homogeneous path is a sharing path for which the edge homogeneity of each edge is positive, i.e., a sharing path composed only of homogeneous links.

Results and Discussion

Anatomy of Cascades. We begin our analysis by characterizing the statistical signature of cascades as they relate to information type. We analyze the three types—science news, conspiracy rumors, and trolling—and find that size and maximum degree are power-law distributed for all three categories. The maximum cascade size values are 952 for science news, 2,422 for conspiracy news, and 3,945 for trolling, and the estimated exponents γ for the power-law distributions are 2.21 for science news, 2.47 for conspiracy, and 2.44 for trolling posts. Tree height values range from 1 to 5, with a maximum height of 5 for science news and conspiracy theories and a maximum height of 4 for trolling. The resulting network is very dense. Notice that such a feature weakens the role of hubs in rumor-spreading dynamics. For further information see *SI Appendix*, section 2.1.

Fig. 1 shows the probability density function (PDF) of the cascade lifetime (using hours as time units) for science and conspiracy. We compute the lifetime as the length of time between the first user and the last user sharing a post. In both categories we find a first peak at ∼1–2 h and a second at ∼20 h, indicating that the temporal sharing patterns are similar irrespective of the difference in topic. We also find that a significant percentage of the information diffuses rapidly (24.42% of the science news and 20.76% of the conspiracy rumors diffuse in less than 2 h, and 39.45% of science news and 40.78% of conspiracy theories in less than 5 h). Only 26.82% of the diffusion of science news and 17.79% of conspiracy lasts more than 1 d.

In Fig. 2 we show the lifetime as a function of the cascade size. For science news we have a peak in the lifetime corresponding to a cascade size value of ≈ 200, and higher cascade size values correspond to high lifetime variability. For conspiracy-related content the lifetime increases with cascade size.

These results suggest that news assimilation differs according to the categories. Science news is usually assimilated, i.e., it reaches a higher level of diffusion quickly, and a longer lifetime does not correspond to a higher level of interest. Conversely, conspiracy rumors are assimilated more slowly and show a positive relation between lifetime and size. For both science and conspiracy news, we compute the size as a function of the lifetime and confirm that differentiation in the sharing patterns is content-driven, and that for conspiracy there is a positive relation between size and lifetime (see *SI Appendix*, section 2.1 for further details).

Homogeneous Clusters. We next examine the social determinants that drive sharing patterns and we focus on the role of homogeneity in friendship networks.

Fig. 3 shows the PDF of the mean-edge homogeneity, computed for all cascades of science news and conspiracy theories. It shows that the majority of links between consecutively sharing users is homogeneous. In particular, the average edge homogeneity value of the entire sharing cascade is always greater than or equal to zero, indicating that either the information transmission occurs inside homogeneous clusters in which all links are homogeneous or it occurs inside mixed neighborhoods in which the balance between homogeneous and nonhomogeneous links is favorable toward the former ones. However, the probability of close to zero mean-edge homogeneity is quite small. Contents tend to circulate only inside the echo chamber.

Hence, to further characterize the role of homogeneity in shaping sharing cascades, we compute cascade size as a function of mean-edge homogeneity for both science and conspiracy news (Fig. 4). In science news, higher levels of mean-edge homogeneity in the interval (0.5, 0.8) correspond to larger cascades, but in conspiracy theories lower levels of mean-edge homogeneity (~ 0.25) correspond to larger cascades. Notice that, although viral patterns related to distinct contents differ, homogeneity is clearly the driver of information diffusion. In other words, different contents generate different echo chambers, characterized by a high level of homogeneity inside them. The PDF of the edge homogeneity, computed for science and conspiracy news as well as the two taken together—both in the unconditional case and in the conditional case (in the event that the user that made the first share in the couple has a positive or negative polarization)—confirms the roughly null probability of a negative edge homogeneity (*SI Appendix*, section 2.1).

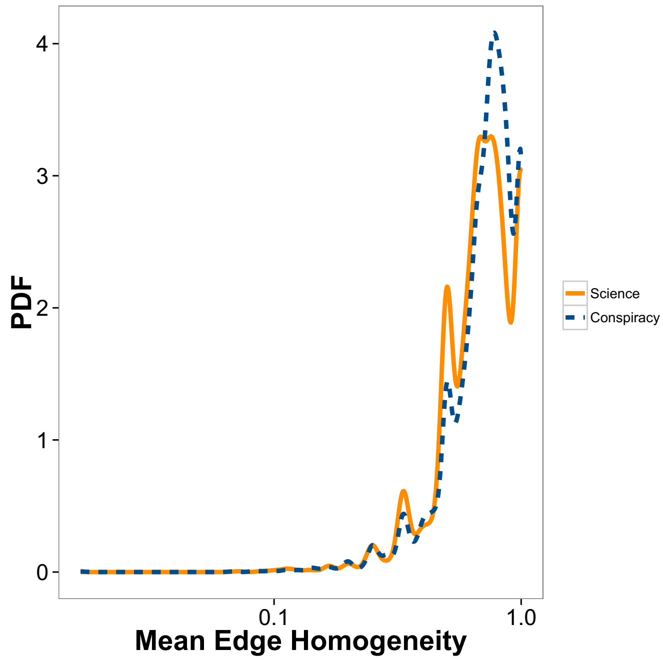

Fig. 3. PDF of edge homogeneity for science (orange) and conspiracy (blue) news. Homogeneity paths are dominant on the whole cascades for both scientific and conspiracy news.

We record the complementary cumulative distribution function (CCDF) of the number of all sharing paths* on each tree compared with the CCDF of the number of homogeneous paths for science and conspiracy news, and the two together. A Kolmogorov–Smirnov test and Q-Q plots confirm that for all three pairs of distributions considered there is no significant statistical difference (see *SI Appendix*, section 2.2 for more details). We confirm the pervasiveness of homogeneous paths.

Indeed, cascades' lifetimes of science and conspiracy news exhibit a probability peak in the first 2 h, and then in the following hours they rapidly decrease. Despite the similar consumption patterns, cascade lifetime expressed as a function of the cascade size differs greatly for the different content sets. However, homogeneity remains the main driver of cascades' propagation. The distributions of the number of total and homogeneous sharing paths are very similar for both content categories. Viral patterns related to contents belonging to different narratives differ, but homogeneity is the primary driver of content diffusion.

*Recall that a sharing path is here defined as any path from the root to one of the leaves of the sharing tree. A homogeneous path is a sharing path for which the edge homogeneity of each edge is positive.

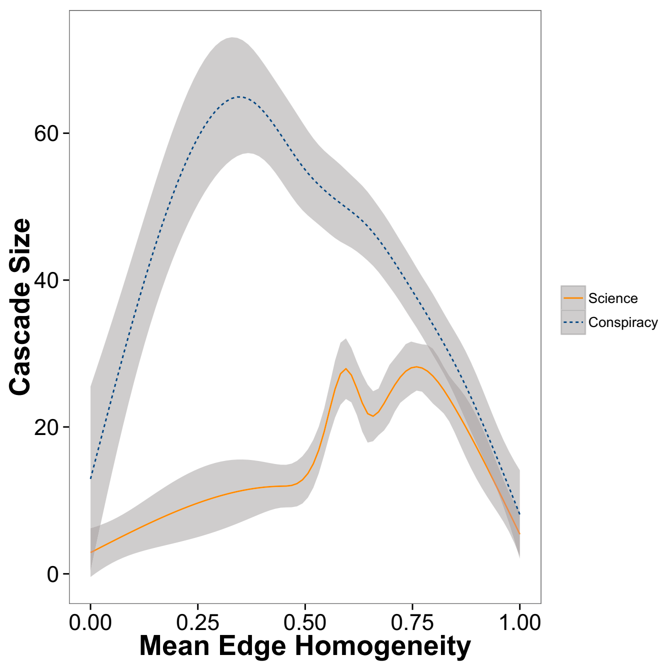

Fig. 4. Cascade size as a function of edge homogeneity for science (orange) and conspiracy (dashed blue) news.

The Model. Our findings show that users mostly tend to select and share content according to a specific narrative and to ignore the rest. This suggests that the determinant for the formation of echo chambers is confirmation bias. To model this mechanism we now introduce a percolation model of rumor spreading to account for homogeneity and polarization. We consider n users connected by a small-world network (41) with rewiring probability r. Every node has an opinion ω_i, $i \in \{1, n\}$ uniformly distributed between [0,1] and is exposed to m news items with a content ϑ_j, $j \in \{1, m\}$ uniformly distributed in [0,1]. At each step the news items are diffused and initially shared by a group of first sharers. After the first step, the news recursively passes to the neighborhoods of previous step sharers, e.g., those of the first sharers during the second step. If a friend of the previous step sharers has an opinion close to the fitness of the news, then she shares the news again.

When

$$\left|\omega_i - \vartheta_j\right| \le \delta,$$

user i shares news j; δ is the sharing threshold.

Because δ by itself cannot capture the homogeneous clusters observed in the data, we model the connectivity pattern as a signed network (4, 42) considering different fractions of homogeneous links and hence restricting diffusion of news only to homogeneous links. We define ϕ_{HL} as the fraction of homogeneous links in the network, M as the number of total links, and n_h as the number of homogeneous links; thus, we have

$$\phi_{HL} = \frac{n_h}{M}, \; 0 \le n_h \le M.$$

Notice that $0 \le \phi_{HL} \le 1$ and that $1 - \phi_{HL}$, the fraction of nonhomogeneous links, is complementary to ϕ_{HL}. In particular, we can reduce the parameters space to $\phi_{HL} \in [0.5, 1]$ as we would restrict our attention to either one of the two complementary clusters.

The model can be seen as a branching process where the sharing threshold δ and neighborhood dimension z are the key parameters. More formally, let the fitness θ_j of the jth news and the opinion ω_i of a the ith user be uniformly independent identically distributed (i.i.d.) between [0,1]. Then the probability p that a user i shares a post j is defined by a probability $p = \min(1, \theta + \delta) - \max(0, \theta - \delta) \approx 2\delta$, because θ and ω are uniformly i.i.d. In general, if ω and θ have distributions $f(\omega)$ and $f(\theta)$, then p will depend on θ,

$$p_\theta = f(\theta) \int_{\max(0, \theta-\delta)}^{\min(1, \theta+\delta)} f(\omega) d\omega.$$

If we are on a tree of degree z (or on a sparse lattice of degree $z+1$), the average number of sharers (the branching ratio) is defined by

$$\mu = zp \approx 2\delta \, z,$$

with a critical cascade size $S = (1-\mu)^{-1}$. If we assume that the distribution of the number m of the first sharers is $f(m)$, then the average cascade size is

$$S = \sum_m f(m) m (1-\mu)^{-1} = \frac{\langle m \rangle_f}{1-\mu} \approx \frac{\langle m \rangle_f}{1 - 2\delta z},$$

where $\langle \ldots \rangle_f = \sum_m \ldots f(m)$ is the average with respect to f. In the simulations we fixed neighborhood dimension $z = 8$ because the branching ratio μ depends upon the product of z and δ and, without loss of generality, we can consider the variation of just one of them.

If we allow a probability q that a neighbor of a user has a different polarization, then the branching ratio becomes $\mu = z(1-q)p$. If a lattice has a degree distribution $d(k)$ ($k = z+1$), we can then assume a usual percolation process that provides a critical branching ratio and that is linear in $\langle k^2 \rangle_d / \langle k \rangle_d$ ($\mu \approx (1-q)p\langle z^2 \rangle / \langle z \rangle$).

Simulation Results. We explore the model parameters space using $n = 5{,}000$ nodes and $m = 1{,}000$ news items with the number of first sharers distributed as (*i*) inverse Gaussian, (*ii*) log normal, (*iii*) Poisson, (*iv*) uniform distribution, and as the real-data distribution (from the science and conspiracy news sample). In Table 1 we show a summary of relevant statistics (min value, first quantile, median, mean, third quantile, and max value) to compare the real-data first sharers distribution with the fitted distributions.[†]

Along with the first sharers distribution, we vary the sharing threshold δ in the interval $[0.01, 0.05]$ and the fraction of homogeneous links ϕ_{HL} in the interval $[0.5, 1]$. To avoid biases induced by statistical fluctuations in the stochastic process, each point of the parameter space is averaged over 100 iterations. $\phi_{HL} \sim 0.5$ provides a good estimate of real-data values. In particular, consistently with the division of in two echo chambers (science and conspiracy), the network is divided into two clusters in which news items remain inside and are transmitted solely within each community's echo chamber (see *SI Appendix*, section 3.2 for the details of the simulation results).

In addition to the science and conspiracy content sharing trees, we downloaded a set of 1,072 sharing trees of intentionally false information from troll pages. Frequently troll information, e.g., parodies of conspiracy theories such as chem-trails containing the active principle of Viagra, is picked up and shared by habitual conspiracy theory consumers. We computed the mean and SD of size and height of all trolling sharing trees, and reproduced the data using our model.[‡] We used fixed parameters from trolling messages

[†]For details on the parameters of the fitted distributions used, see *SI Appendix*, section 3.2.

[‡]Note that the real-data values for the mean (and SD) of size and height on the troll posts are, respectively, 23.54 (122.32) and 1.78 (0.73).

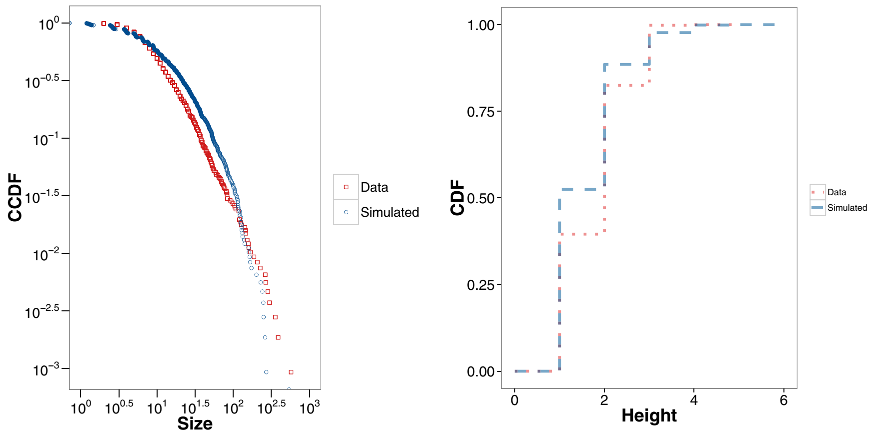

Fig. 5. CCDF of size (*Left*) and CDF of height (*Right*) for the best parameters combination that fits real-data values, $(\phi_{HL}, r, \delta) = (0.56, 0.01, 0.015)$, and first sharers distributed as $IG(18.73, 9.63)$.

sample (the number of nodes in the system and the number of news items) and varied the fraction of homogeneous links ϕ_{HL}, the rewiring probability r, and sharing threshold δ. See *SI Appendix*, section 3.2 for the distribution of first sharers used and for additional simulation results of the fit on trolling messages.

We simulated the model dynamics with the best combination of parameters obtained from the simulations and the number of first sharers distributed as an inverse Gaussian. Fig. 5 shows the CCDF of cascades' size and the cumulative distribution function (CDF) of their height. A summary of relevant statistics (min value, first quantile, median, mean, third quantile, and max value) to compare the real-data size and height distributions with the fitted ones is reported in *SI Appendix*, section 3.2.

We find that the inverse Gaussian is the distribution that best fits the data both for science and conspiracy news, and for troll messages. For this reason, we performed one more simulation using the inverse Gaussian as distribution of the number of first sharers, 1,072 news items, 16,889 users, and the best parameters combination obtained in the simulations.[§] The CCDF of size and the CDF of height for the above parameters combination, as well as basic statistics considered, fit real data well.

Conclusions

Digital misinformation has become so pervasive in online social media that it has been listed by the WEF as one of the main threats to human society. Whether a news item, either substantiated or not, is accepted as true by a user may be strongly affected by social norms or by how much it coheres with the user's system of beliefs (32, 33). Many mechanisms cause false information to gain acceptance, which in turn generate false beliefs that, once adopted by an individual, are highly resistant to correction (34–37). In this work, using extensive quantitative analysis and data-driven modeling, we provide important insights toward the understanding of the mechanism behind rumor spreading. Our findings show that users mostly tend to select and share content related to a specific narrative and to ignore the rest. In particular, we show that social homogeneity is the primary driver of content diffusion, and one frequent result is the formation of homogeneous, polarized clusters. Most of the times the information is taken by a friend having the same profile (polarization)—i.e., belonging to the same echo chamber.

We also find that although consumers of science news and conspiracy theories show similar consumption patterns with respect to content, their cascades differ.

Our analysis shows that for science and conspiracy news a cascade's lifetime has a probability peak in the first 2 h, followed by a rapid decrease. Although the consumption patterns are similar, cascade lifetime as a function of the size differs greatly.

These results suggest that news assimilation differs according to the categories. Science news is usually assimilated, i.e., it reaches a higher level of diffusion, quickly, and a longer lifetime does not correspond to a higher level of interest. Conversely, conspiracy rumors are assimilated more slowly and show a positive relation between lifetime and size.

The PDF of the mean-edge homogeneity indicates that homogeneity is present in the linking step of sharing cascades. The distributions of the number of total sharing paths and homogeneous sharing paths are similar in both content categories.

Viral patterns related to distinct contents are different but homogeneity drives content diffusion. To mimic these dynamics, we introduce a simple data-driven percolation model of signed networks, i.e., networks composed of signed edges accounting for nodes preferences toward specific contents. Our model reproduces the observed dynamics with high accuracy.

Users tend to aggregate in communities of interest, which causes reinforcement and fosters confirmation bias, segregation, and polarization. This comes at the expense of the quality of the information and leads to proliferation of biased narratives fomented by unsubstantiated rumors, mistrust, and paranoia.

According to these settings algorithmic solutions do not seem to be the best options in breaking such a symmetry. Next envisioned steps of our research are to study efficient communication strategies accounting for social and cognitive determinants behind massive digital misinformation.

Table 1. Summary of relevant statistics comparing synthetic data with the real ones

Values	Data	*IG*	Lognormal	Poisson
Min	1	0.36	0.10	20
First quantile	5	4.16	3.16	35
Median	10	10.45	6.99	39
Mean	39.34	39.28	13.04	39.24
Third quantile	27	31.59	14.85	43
Max	3,033	1814	486.10	66

The inverse Gaussian (*IG*) shows the best fit for the distribution of first sharers with respect to all of the considered statistics.

[§]The best parameters combinations is $\phi_{HL} = 0.56$, $r = 0.01$, $\delta = 0.015$. In this case we have a mean size equal to 23.42 (33.43) and a mean height 1.28 (0.88), and it is indeed a good approximation; see *SI Appendix*, section 3.2.

ACKNOWLEDGMENTS. Special thanks go to Delia Mocanu, "Protesi di Protesi di Complotto," "Che vuol dire reale," "La menzogna diventa verita e passa alla storia," "Simply Humans," "Semplicemente me," Salvatore Previti, Elio Gabalo, Sandro Forgione, Francesco Pertini, and "The rooster on the trash" for their valuable suggestions and discussions. Funding for this work was provided by the EU FET Project MULTIPLEX, 317532, SIMPOL, 610704, the FET Project DOLFINS 640772, SoBigData 654024, and CoeGSS 676547.

1. Brown J, Broderick AJ, Lee N (2007) Word of mouth communication within online communities: Conceptualizing the online social network. *J Interact Market* 21(3):2–20.
2. Kahn R, Kellner D (2004) New media and internet activism: From the "battle of Seattle" to blogging. *New Media Soc* 6(1):87–95.
3. Quattrociocchi W, Conte R, Lodi E (2011) Opinions manipulation: Media, power and gossip. *Adv Complex Syst* 14(4):567–586.
4. Quattrociocchi W, Caldarelli G, Scala A (2014) Opinion dynamics on interacting networks: Media competition and social influence. *Sci Rep* 4:4938.
5. Kumar R, Mahdian M, McGlohon M (2010) Dynamics of conversations. *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (ACM, New York), pp 553–562.
6. Sunstein C, Vermeule A (2009) Conspiracy theories: Causes and cures. *J Polit Philos* 17(2):202–227.
7. Kadlec C (2011) The goal is power: The global warming conspiracy. *Forbes*, July 25, 2011. Available at www.forbes.com/sites/charleskadlec/2011/07/25/the-goal-is-po.wer-the-global-warming-conspiracy/. Accessed August 21, 2015.
8. Millman J (2014) The inevitable rise of Ebola conspiracy theories. *The Washington Post*, Oct. 13, 2014. Available at https://www.washingtonpost.com/news/wonk/wp/2014/10/13/the-inevitable-rise-of-ebola-conspiracy-theories/. Accessed August 31, 2015.
9. Lamothe D (2015) Remember Jade Helm 15, the controversial military exercise? It's over. *The Washington Post*, Sept. 14, 2015. Available at https://www.washingtonpost.com/news/checkpoint/wp/2015/09/14/remember-jade-helm-15-the-controversial-military-exercise-its-over/. Accessed September 20, 2015.
10. Bessi A, et al. (2015) Science vs conspiracy: Collective narratives in the age of misinformation. *PLoS One* 10(2):e0118093.
11. Mocanu D, Rossi L, Zhang Q, Karsai M, Quattrociocchi W (2015) Collective attention in the age of (mis) information. *Comput Human Behav* 51:1198–1204.
12. Bessi A, Scala A, Rossi L, Zhang Q, Quattrociocchi W (2014) The economy of attention in the age of (mis) information. *J Trust Manage* 1(1):1–13.
13. Furedi F (2006) *Culture of Fear Revisited* (Bloomsbury, London).
14. Aiello LM, et al. (2012) Friendship prediction and homophily in social media. *ACM Trans Web* 6(2):9.
15. Gu B, Konana P, Raghunathan R, Chen HM (2014) Research note—the allure of homophily in social media: Evidence from investor responses on virtual communities. *Inf Syst Res* 25(3):604–617.
16. Bessi A, et al. (2015) Viral misinformation: The role of homophily and polarization. *Proceedings of the 24th International Conference on World Wide Web Companion* (International World Wide Web Conferences Steering Committee, Florence, Italy), pp 355–356.
17. Bessi A, et al. (2015) Trend of narratives in the age of misinformation. *PLoS One* 10(8): e0134641.
18. Zollo F, et al. (2015) Emotional dynamics in the age of misinformation. *PLoS One* 10(9):e0138740.
19. Byford J (2011) *Conspiracy Theories: A Critical Introduction* (Palgrave Macmillan, London).
20. Fine GA, Campion-Vincent V, Heath C (2005) Rumor Mills: The Social Impact of Rumor and Legend, eds Fine GA, Campion-Vincent V, Heath C (Aldine Transaction, New Brunswick, NJ), pp 103–122.
21. Hogg MA, Blaylock DL (2011) *Extremism and the Psychology of Uncertainty* (John Wiley & Sons, Chichester, UK), Vol 8.
22. Betsch C, Sachse K (2013) Debunking vaccination myths: Strong risk negations can increase perceived vaccination risks. *Health Psychol* 32(2):146–155.
23. Howell L (2013) Digital wildfires in a hyperconnected world. WEF Report 2013. Available at reports.weforum.org/global-risks-2013/risk-case-1/digital-wildfires-in-a-hyperconnected-world. Accessed August 31, 2015.
24. Qazvinian V, Rosengren E, Radev DR, Mei Q (2011) Rumor has it: Identifying misinformation in microblogs. *Proceedings of the Conference on Empirical Methods in Natural Language Processing* (Association for Computational Linguistics, Stroudsburg, PA), pp 1589–1599.
25. Ciampaglia GL, et al. (2015) Computational fact checking from knowledge networks. arXiv:1501.03471.
26. Resnick P, Carton S, Park S, Shen Y, Zeffer N (2014) Rumorlens: A system for analyzing the impact of rumors and corrections in social media. *Proceedings of Computational Journalism Conference* (ACM, New York).
27. Gupta A, Kumaraguru P, Castillo C, Meier P (2014) Tweetcred: Real-time credibility assessment of content on twitter. *Social Informatics* (Springer, Berlin), pp 228–243.
28. Al Mansour AA, Brankovic L, Iliopoulos CS (2014) A model for recalibrating credibility in different contexts and languages-a twitter case study. *Int J Digital Inf Wireless Commun* 4(1):53–62.
29. Ratkiewicz J, et al. (2011) Detecting and tracking political abuse in social media. *Proceedings of the 5th International AAAI Conference on Weblogs and Social Media* (AAAI, Palo Alto, CA).
30. Dong XL, et al. (2015) Knowledge-based trust: Estimating the trustworthiness of web sources. *Proc VLDB Endowment* 8(9):938–949.
31. Nyhan B, Reifler J, Richey S, Freed GL (2014) Effective messages in vaccine promotion: A randomized trial. *Pediatrics* 133(4):e835–e842.
32. Zhu B, et al. (2010) Individual differences in false memory from misinformation: Personality characteristics and their interactions with cognitive abilities. *Pers Individ Dif* 48(8):889–894.
33. Frenda SJ, Nichols RM, Loftus EF (2011) Current issues and advances in misinformation research. *Curr Dir Psychol Sci* 20(1):20–23.
34. Kelly GR, Weeks BE (2013) The promise and peril of real-time corrections to political misperceptions. *Proceedings of the 2013 Conference on Computer Supported Cooperative Work* (ACM, New York), pp 1047–1058.
35. Meade ML, Roediger HL, 3rd (2002) Explorations in the social contagion of memory. *Mem Cognit* 30(7):995–1009.
36. Koriat A, Goldsmith M, Pansky A (2000) Toward a psychology of memory accuracy. *Annu Rev Psychol* 51(1):481–537.
37. Ayers MS, Reder LM (1998) A theoretical review of the misinformation effect: Predictions from an activation-based memory model. *Psychon Bull Rev* 5(1):1–21.
38. Sunstein C (2001) *Echo Chambers* (Princeton Univ Press, Princeton, NJ).
39. Kelly GR (2009) Echo chambers online?: Politically motivated selective exposure among internet news users. *J Comput Mediat Commun* 14(2):265–285.
40. Facebook. (2015) Using the graph API. Available at https://developers.facebook.com/docs/graph-api/using-graph-api. Accessed December 19, 2015.
41. Watts DJ, Strogatz SH (1998) Collective dynamics of 'small-world' networks. *Nature* 393(6684):440–442.
42. Leskovec J, Huttenlocher D, Kleinberg J (2010) Signed networks in social media. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (ACM, New York), pp 1361–1370.