

# Homophily and polarization in the age of misinformation

Alessandro Bessi<sup>1,2</sup>, Fabio Petroni<sup>3</sup>, Michela Del Vicario<sup>2</sup>, Fabiana Zollo<sup>2</sup>, Aris Anagnostopoulos<sup>3</sup>, Antonio Scala<sup>4</sup>, Guido Caldarelli<sup>2</sup>, and Walter Quattrociocchi<sup>1,a</sup>

<sup>1</sup> IUSS, Pavia, Italy

<sup>2</sup> IMT Institute for Advanced Studies, Lucca, Italy

<sup>3</sup> Sapienza University, Roma, Italy

<sup>4</sup> ISC-CNR, Roma, Italy

Received 1 December 2015 / Received in final form 30 January 2016  
Published online 26 October 2016

**Abstract.** The World Economic Forum listed massive digital misinformation as one of the main threats for our society. The spreading of unsubstantiated rumors may have serious consequences on public opinion such as in the case of rumors about Ebola causing disruption to health-care workers. In this work we target Facebook to characterize information consumption patterns of 1.2M Italian users with respect to verified (science news) and unverified (conspiracy news) contents. Through a thorough quantitative analysis we provide important insights about the anatomy of the system across which misinformation might spread. In particular, we show that users' engagement on verified (or unverified) content correlates with the number of friends having similar consumption patterns (*homophily*). Finally, we measure how this social system responded to the injection of 4,709 false information. We find that the frequent (and selective) exposure to specific kind of content (*polarization*) is a good proxy for the detection of homophile clusters where certain kind of rumors are more likely to spread.

## 1 Introduction

The Web has become pervasive and digital technology permeates every aspect of daily life. Social interaction, healthcare activity, political engagement, and economic decision-making are influenced by the digital hyper-connectivity [1–8]. Nowadays, everyone can produce and access a variety of information by actively participating in the diffusion and reinforcement of narratives. Such a shift of paradigm in information consumption has profoundly affected the way users get informed [9–15]. The spreading of unsubstantiated rumors, whether intentional or unintentional, could have serious consequences; the World Economic Forum has listed *massive digital misinformation* as one of the main risks for the modern society [16]. Interesting is the popular case of Senator Cirenga's [17, 18] law, proposing to fund parliamentary deputy members with 134 million of euros (10% of the Italian GDP) in case of defeat in the political

<sup>a</sup> e-mail: [walterquattrociocchi@gmail.com](mailto:walterquattrociocchi@gmail.com)

competition. This was an intentional joke – the text of the post was explicitly mentioning its provocative nature – which became popular within online political activists to an extent that it has been used as an argumentation in political debates [19]. In this work we focus on two distinct types of news – science and conspiracy – which differ for the possibility of verifying their content. On the one hand, science news have pointers to the related scientific papers as well as references to authors, universities, institutions and recognized organizations. On the other hand, conspiracy news diffuse information that are “neglected” by main stream media; thus it is almost impossible to verify their sources. Indeed, science news are aiming at diffusing scientific knowledge and scientific thinking, whereas conspiracy news provide alternative arguments that are difficult to document. Conspiracists tend to reduce the complexity of reality by explaining significant social or political events as secret plots conceived by powerful individuals or organizations. Just recently, the fear of an Ebola outbreak in the United States rippled through social media networks [20–22]. Furthermore misinformation may be particularly difficult to correct [23–26]. In fact, it has been recently shown [27] that conspiracy and mainstream information reverberate in a similar way on social media and that users generally exposed to conspiracy stories are more prone to like and share satirical information [28]. In this work we analyze a sample of 1.2 M Facebook Italian users consuming scientific and conspiracy news. Our findings show that users’ engagement on a specific content correlates with the number of friends having similar consumption patterns (*homophily*). We then test the relationship between the usual exposure (polarization) to undocumented rumors (conspiracy stories) with respect to the permeability to deliberate false information – 4,709 intentional satirical false claims. Our work provides important insights about the understanding of the diffusion of unverified rumors. In particular, we show that through polarization, we can detect homophily clusters where misleading rumors are more likely to spread. Conversely, such social patterns might represent serious warnings about the effects of current algorithmic specifications for content provisioning. Indeed, the tendency to aggregate in homophile clusters might foster social reinforcement and confirmation bias and thus polarization [25, 26]. Along this path, recent studies showed that the primary driver for viral phenomena is confirmation bias and that cascades are bounded to the echo chamber size [29].

## 2 Related work

**Rumor spreading.** A wide branch of the literature is devoted to understanding the spread of rumors and behaviors on online social networks, focusing both on their structural properties and on their effects on social dynamics. In [10] the author investigates the effects of the topology on diffusion, showing that the network structure has a significant effect. Moreover, he observes that rumors and behaviors spread farther and faster across clustered-lattice networks than across random networks. In [30] authors find that the probability of contagion is tightly controlled by the number of connected components in an individual’s contact neighborhood, rather than by the actual size of the neighborhood. In [31] researchers show that although long ties are relevant for spreading information about an innovation or social movement, they are not sufficient w.r.t. the social reinforcement necessary to act on that information. In [32] a method for identifying influence and susceptibility in networks is presented. Their estimations show that influential individuals with influential friends may allow the spread of information in the network. Moreover, a key factor in identifying true contagion in social network is to distinguish between peer-to-peer influence and homophily: in the first case, a node influences or causes outcomes to its neighbors, whereas in the second one dyadic similarities between nodes create correlated

outcome patterns among neighbors that could mimic viral contagions even without direct causal influence [33]. The study presented in [34] reveals that there is a substantial level of topical similarity among users who are close to each other in the social network, suggesting that users with similar interests are more likely to be friends. In [35] authors have developed an estimation framework to distinguish influence and homophily effects in dynamic networks and find that homophily explains more than 50% of the perceived behavioral contagion. In [36] the role of social networks and exposure to friends' activities in information resharing on Facebook is analyzed. Once having isolated contagion from other confounding effects such as homophily, authors claim that there is a considerably higher chance to share contents when users are exposed to friends' resharing. Network homophily has been also introduced into the dynamics of cultural interaction, leading to a model in which patterns of social interaction change with processes of social influence [37].

**Cascades.** Recent studies aim at unfolding cascades (i.e., a consecutive series of re-shares) characteristics. A large variety of settings has been explored, such as blogging [38–40], e-mail [41, 42], and social network services – e.g., Twitter [43–45]. Some works [46, 47] also consider the anatomy and the predictivity of large Facebook cascades, showing that in general rumor cascades run deeper in the social network than reshare cascades [48]. In particular, in [46] authors find a method for predicting whether a cascade will continue to grow in the future. Also temporal and structural features are predictors of cascades size.

**Diffusion of unsubstantiated rumors.** In the last years, a new online phenomenon is emerging i.e., the spreading of unsubstantiated and false claims through online social networks. It has been shown that online unsubstantiated rumors (e.g., the link between vaccines and autism, the global warming induced by chem-trails, the secret alien government) reverberate in a comparable way with respect to mainstream information such as scientific news and updates [27]. Recent works [27, 28] reveal that massive digital misinformation permeates online social dynamics creating viral phenomena even on intentional parodistic false information. In [26] authors show that consumption patterns around information supporting different (and opposite) world views are similar.

### 3 Data Collection

We identified two main categories of pages: conspiracy news – i.e., pages promoting contents *neglected* by main stream media – and science news. We defined the space of our investigation with the help of Facebook groups very active in debunking conspiracy theses (“*Protesi di Protesi di Complotto*”, “*Che vuol dire reale*”, “*La menzogna diventa verità e passa alla storia*”). As an additional control, we used the self-description of a page to determine its focus.

**Conspiracy and Science news.** The resulting dataset contains 73 public Facebook pages; 34 of such pages are related to scientific news while the other 39 to news that may be considered conspiratorial; we refer to the former as *science pages* and to the latter as *conspiracy pages*. Notice that the dataset used in the analysis is the same used in [26] and [28]. Table 1 summarizes the details of our data collection. We downloaded all the posts from these pages in a timespan of 4 years (2010 to 2014). In addition, we collected all the *likes* and *comments* from the posts, and we counted the number of *shares*. In total, we collected around 9 M likes and 1 M comments, relative to  $\sim 1.2$  M and  $\sim 280$  K Facebook users, respectively (see Table 1). Likes, shares, and comments have a different meaning from the user viewpoint. Most of the times, a like stands for a positive feedback to the post; a share expresses the will to increase the

**Table 1.** Breakdown of Facebook dataset. The number of pages, posts, likes, comments, and shares for science and conspiracy pages.

	<b>Total</b>	<b>Science</b>	<b>Conspiracy</b>
Pages	73	34	39
Posts	271, 296	62, 705	208, 591
Likes	9, 164, 781	2, 505, 399	6, 659, 382
Comments	1, 017, 509	180, 918	836, 591
Shares	17, 797, 819	1, 471, 088	16, 326, 731
Likers	1, 196, 404	332, 357	864, 047
Commenters	279, 972	53, 438	226, 534

visibility of a given information; and a comment is the way in which online collective debates take form. Comments may contain negative or positive feedbacks with respect to the post.

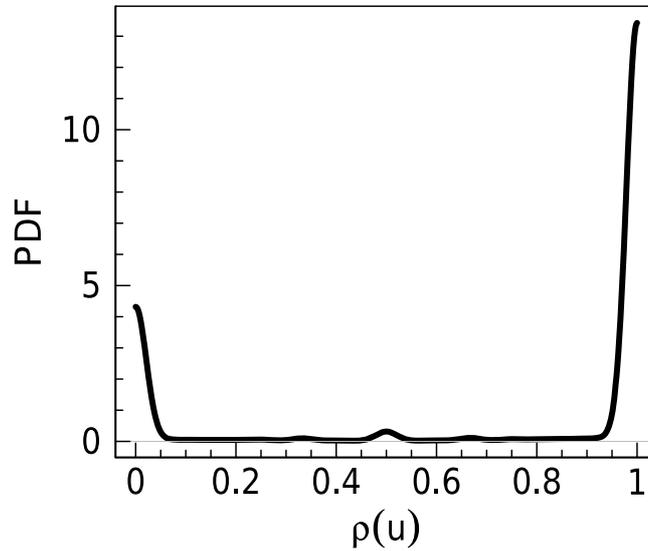
**Ego Networks.** In addition, we collected the ego networks of users who liked at least one post on science or conspiracy pages<sup>1</sup>.

**Troll Pages.** As a control, we used 4,709 posts from two satirical Facebook pages (which we will refer to as *troll posts* and *troll pages*) promoting intentionally false and caricatural version of the most debated issues. The most popular of the two is called “Semplicemente Me” (10 K followers) [49] and it is focused on general online rumors. The second one is called “Simply Humans” (1 K followers) [50], and mostly hosts posts of conspiratorial nature. We collected about 40 K likes and 59 K comments on these posts, performed by about 16 K and 43 K Facebook users, respectively. The contents of these pages were able to trigger several viral phenomena, with one of them reaching more than 100 K shares. We use troll memes as a benchmark to test users interaction with information that are deliberately false.

## 4 Preliminaries and definitions

In this section we provide some of the basic definitions and notions that we use throughout the paper. Let  $\mathcal{P}$  be the set of all the posts in our collection, and  $\mathcal{P}_{\text{science}}$  ( $\mathcal{P}_{\text{consp}}$ ) be the set of posts of the 34 (39) Facebook pages about science (conspiracy) news. Let  $V$  be the set of all the 1.2 M users and  $E$  the edges representing their Facebook friendship connections; these sets define a graph  $G = (V, E)$ . Hence, the graph of likes on a post,  $G^{\mathbf{L}} = (V^{\mathbf{L}}, E^{\mathbf{L}})$  is the subgraph of  $G$  whose users have liked a post. Thus,  $V^{\mathbf{L}}$  is the set of users of  $V$  who have liked at least one post, and we set  $E^{\mathbf{L}} = \{(u, v) \in E; u, v \in V^{\mathbf{L}}\}$ . Following previous works [26–28], we study the polarization of users i.e., the tendency of users to interact with only a single type of information; in particular, we study the polarization towards science and conspiracy. Formally we define the *polarization*  $\rho(u) \in [0, 1]$  of user  $u \in V^{\mathbf{L}}$  as the ratio of likes that  $u$  has performed on conspiracy posts: assuming that  $u$  has performed  $x$  and  $y$  likes on science and conspiracy posts, respectively, we let  $\rho(u) = y/(x + y)$ . Thus, a user  $u$  for whom  $\rho(u) = 0$  is polarized towards science, whereas a user with  $\rho(u) = 1$  is polarized towards conspiracy. Note that we ignore the commenting activity since a comment may be an endorsement, a criticism, or even a response to a previous comment. Furthermore, we define the *engagement*  $\psi(u)$  of a user  $u$  as her liking

<sup>1</sup> We used publicly available data, so we collected only data for which the users had the corresponding permissions open.



**Fig. 1.** Polarization on contents. Probability density function (PDF) of users' polarization. Notice the strong bimodality of the distribution, with two sharp peaks localized at  $0 \lesssim \rho \lesssim 0.005$  (science users) and at  $0.95 \lesssim \rho \lesssim 1$  (conspiracy users).

activity normalized with respect to the total number of the users likes in our dataset. By defining  $\theta(u)$  as the total number of likes that the user  $u$  has expressed in posts in our collection  $\mathcal{P}$ , notice that the following condition holds:  $\psi(u) = \frac{\theta(u)}{\max_v \theta(v)}$ .

In Fig. 1 we show that the probability density function (PDF) for the polarization of all the users in  $V^L$  is a sharply peaked bimodal where the vast majority of users are polarized either towards science ( $\rho(u) \sim 0$ ) or conspiracy ( $\rho(u) \sim 1$ ). Hence, Fig. 1 shows that most of likers can be divided into two groups of users, those *polarized towards science* and those *polarized towards conspiracy* news. To better define the properties of these groups, we define the set  $V_{\text{science}}^L$  of users with polarization more than 95% towards science

$$V_{\text{science}}^L = \{u \in V^L; \rho(u) < 0.05\},$$

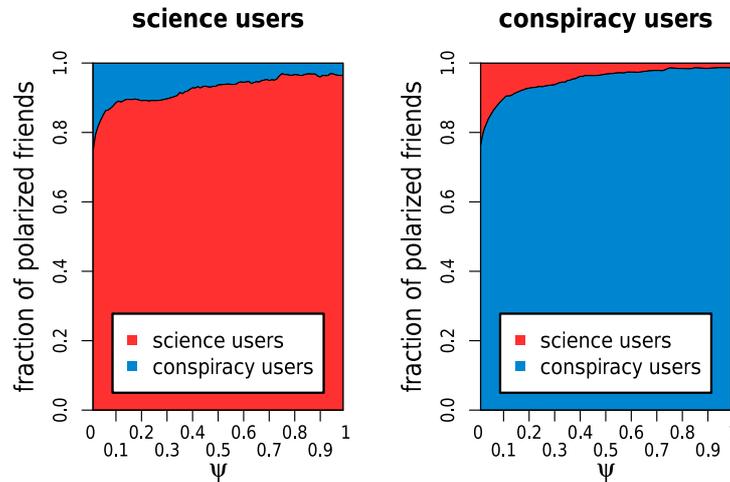
and the set  $V_{\text{consp}}^L$  of users with polarization more than 95% towards conspiracy

$$V_{\text{consp}}^L = \{u \in V^L; \rho(u) > 0.95\};$$

such sets corresponds to the two peaks of the bimodal distribution.

We then define the induced subgraphs of  $G$ ,  $G_{\text{science}}^L = (V_{\text{science}}^L, E_{\text{science}}^L)$  and  $G_{\text{consp}}^L = (V_{\text{consp}}^L, E_{\text{consp}}^L)$  in the natural way, for example,  $E_{\text{science}}^L$  contains all the edges in  $E$  such that both endpoints are in  $V_{\text{science}}^L$ . Such subgraphs describe the internal social network of strongly polarized users.

Finally, we define some of the graph-theoretic terms that we use throughout the paper. The *degree* of node  $u$ ,  $\text{deg}(u)$ , is the number of neighbors of node  $u$ . The  $k$ -*core* of a graph  $H$  is the maximal subgraph  $H'$  of  $H$  such that the degree of each node in  $H'$  is at least  $k$ .



**Fig. 2.** Fraction of polarized neighbors as a function of the engagement  $\psi(\cdot)$ . Left panel: for a polarized scientific user  $u \in V_{\text{science}}^{\text{L}}$ , the fraction of friends  $v$  with the same polarization ( $v \in V_{\text{science}}^{\text{L}}$ ) is very high ( $\gtrsim 0.75$ ) and grows with engagement  $\psi$ . Right panel: for a polarized conspiracy user  $u \in V_{\text{consp}}^{\text{L}}$ , the fraction of friends  $v$  with the same polarization ( $v \in V_{\text{consp}}^{\text{L}}$ ) is very high ( $\gtrsim 0.75$ ) and grows with the engagement  $\psi$ .

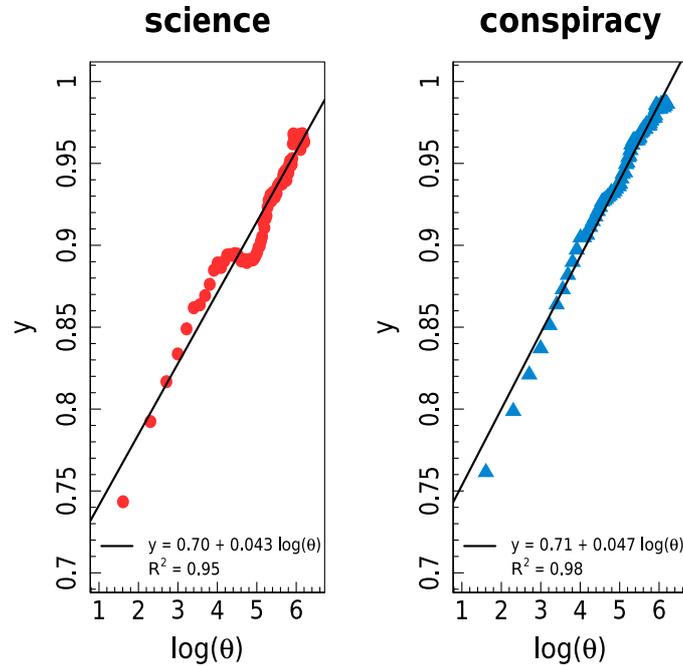
## 5 Results and discussion

### 5.1 Homiphily: Engagement and friends

In this section we want to understand if the engagement in a specific kind of content is a good proxy to detect group of users with similar attitudes. Homophily – i.e., the tendency of users to aggregate around common interests – has been already pointed out as a factor in rumor spreading [34, 51]. Indeed, online content mingle cognitive, behavioral, and social aspects of a user community. The resulting ecosystem allows to investigate the various processes at play in the interactions of individuals, and to study the ways in which users relate with the kind of information they interact with. Thus, users' liking activity across contents of the different categories [26–28] may be intended as the preferential attitude towards the one or the other type of information (documented or not).

**Polarization and the friendship network.** In Fig. 2 we show the fraction of polarized friends of polarized users as a function of their engagement  $\psi(\cdot)$  both in the case of users polarized toward science (left panel) and in the case of users polarized toward conspiracy (right panel). Figure 2 shows that social interactions of Facebook users are driven by homophily: users not only tend to be very polarized, but they also tend to be linked to users with similar preferences. In fact, in both panels of Figure 2 we can observe that, for a polarized scientific (conspiracy) user, the fraction of friends  $v$  with the same polarization is very high ( $\gtrsim 0.75$ ) and grows with the engagement  $\psi$ .

Summarizing, we find that the activity of a user on a specific kind of content increases the probability to have friends with similar characteristics. Such information is a precious insight toward the understanding of information diffusion. Indeed, in [27] we have shown that users usually exposed to undocumented claims (e.g., conspiracy stories) are the the most likely to confuse intentional false information as usual conspiracy stories.



**Fig. 3.** Predicting the number of polarized friends. Left panel: scientific polarized users. Right panel: conspiracy polarized users. In both panels, for a polarized user  $u$ , we plot the fraction of polarized friends with the same polarization (full red circles for scientific, full blue triangles for conspiracy) versus the number of likes  $\log(\theta(u))$  of user  $u$ . Full lines are the results of a linear regression model  $y(u) = \beta_0 + \beta_1 \log(\theta(u))$ . Coefficients are estimated using ordinary least squares; in both cases, all the p-values are close to zero.

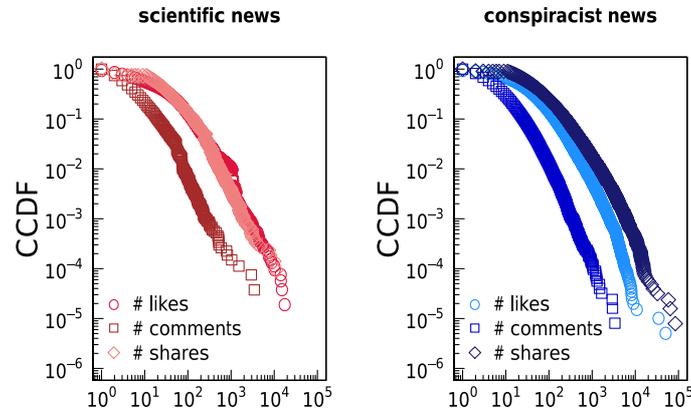
**Predicting homophily.** For a polarized scientific users  $u \in V_{\text{science}}^{\text{L}}$ , in the left panel of Fig. 3, we show the log-linear plot of the average fraction  $y$  friends  $v$  with the same polarization ( $v \in V_{\text{science}}^{\text{L}}$ ) respect given number of likes  $\theta$  of the user  $u$ . In the right panel, we show the same quantities for polarized conspiracy users. Figure 3 suggests in both cases a linear correlation among the variables; thus, we check whether for a polarized user  $u$ , the fraction of polarized friends in its category  $y(u)$  can be predicted by means of a linear regression model where the explanatory variable is a logarithmic transformation of the number of likes  $\theta(u)$ , i.e.

$$y(u) = \beta_0 + \beta_1 \log(\theta(u)).$$

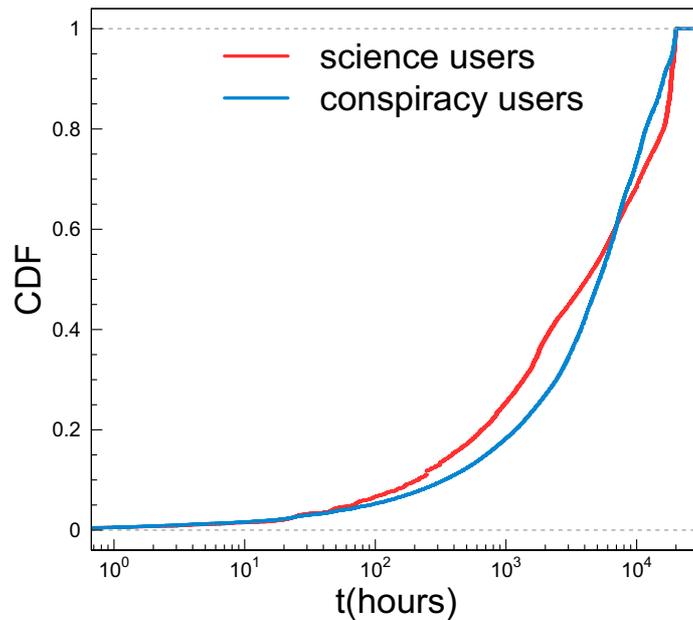
Coefficients are estimated using ordinary least squares and they are – with the corresponding standard errors inside the round brackets –  $\hat{\beta}_0 = 0.70$  (0.005) and  $\hat{\beta}_1 = 0.043$  (0.001), with  $R^2 = 0.95$ , for users polarized towards science, and  $\hat{\beta}_0 = 0.71$  (0.003) and  $\hat{\beta}_1 = 0.047$  (0.0006), with  $R^2 = 0.98$ , for users polarized towards conspiracy. All the p-values are close to zero.

### 5.2 Consumption patterns

In this section we focus on the consumption patterns of polarized users. We first analyze the statistics of likes, comments, and shares of posts in  $\mathcal{P}_{\text{science}}$  and  $\mathcal{P}_{\text{consp}}$ .



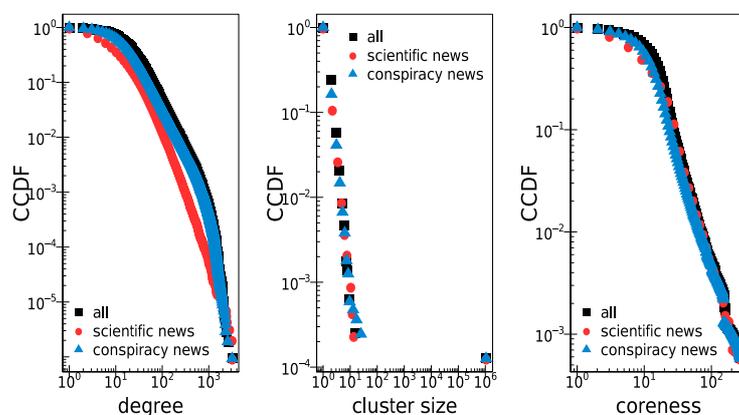
**Fig. 4.** Consumption Patterns. Empirical complementary cumulative distribution function (CCDF) for the number of likes, comments, and shares of  $\mathcal{P}_{\text{science}}$  (left panel) and  $\mathcal{P}_{\text{consp}}$  (right panel).



**Fig. 5.** Users' Lifetime. Empirical cumulative distribution function (CDF) for the temporal distance between the first and last like of users in  $V_{\text{science}}^{\text{L}}$  and  $V_{\text{consp}}^{\text{L}}$  on posts in  $\mathcal{P}_{\text{science}}$  and  $\mathcal{P}_{\text{consp}}$ , respectively; see text for more details.

Figure 4 shows that the empirical complementary cumulative distribution function (CCDF) for all such distributions are heavy-tailed, as often observed in social media analysis; moreover, the consumption patterns both for scientific and for conspiracy users are very similar.

We then analyze the users' persistence on contents. To such an aim, we measure how users in  $V_{\text{science}}^{\text{L}}$  and  $V_{\text{consp}}^{\text{L}}$  keep on consuming information in  $\mathcal{P}_{\text{science}}$  and  $\mathcal{P}_{\text{consp}}$  over time. Figure 5 shows the CDF for the lifetime of  $V_{\text{science}}^{\text{L}}$  and  $V_{\text{consp}}^{\text{L}}$  on posts in  $\mathcal{P}_{\text{science}}$  and  $\mathcal{P}_{\text{consp}}$ , respectively; here the *lifetime* of a user in  $V_{\text{science}}^{\text{L}}$  ( $V_{\text{consp}}^{\text{L}}$ ) is



**Fig. 6.** Network metrics. Complementary cumulative distribution function (CCDF) of the degree of each node, the size of the connected components, and the coreness of each node of the graphs  $G$  (all),  $G_{\text{science}}^{\text{L}}$  (scientific news), and  $G_{\text{consp}}^{\text{L}}$  (conspiracy news). All graphs present similar distributions, hinting that both groups of polarized users have the same structure of social interactions of the whole social network.

the difference of the post time of the last and first post in  $\mathcal{P}_{\text{science}}$  ( $\mathcal{P}_{\text{consp}}$ ) that she liked; we observe that also the persistence on contents is similar for science and for conspiracy users.

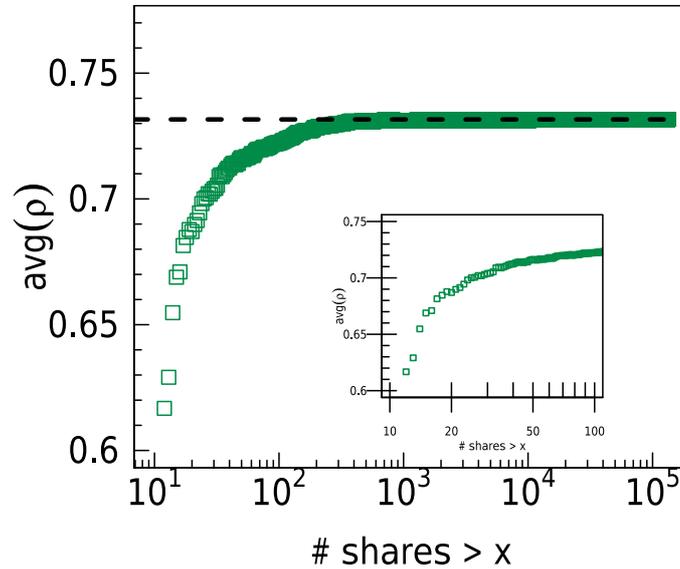
### 5.3 Polarized clusters

In order to understand whether the polarized clusters have different structures, we study the polarized users' network  $G_{\text{science}}^{\text{L}}$  and  $G_{\text{consp}}^{\text{L}}$  compared to the whole social network  $G$ . In Fig. 6, we show the complementary cumulative distribution function (CCDF) of the degree of each node, the size of the connected components, and the coreness of each node for the graphs  $G$ ,  $G_{\text{science}}^{\text{L}}$ , and  $G_{\text{consp}}^{\text{L}}$ . Panels of Fig. 6 show that the social structures  $G_{\text{science}}^{\text{L}}$  and  $G_{\text{consp}}^{\text{L}}$  of polarized users are very similar, both among themselves and to the whole social network  $G$ .

### 5.4 Diffusion of misleading rumors

We have shown that the connectivity among users presents an homophily pattern – i.e, users with similar polarization tend to aggregate together. However, the two groups of polarized users (science and conspiracy) share similar information consumption patterns and internal social network structure. We want now to understand how polarized users react to the inoculation of false information; hence, we need a benchmark sample of posts that can be cleared and unambiguously identified as deliberately false.

**Troll Memes.** We use as a benchmark a sample  $\mathcal{P}_{\text{troll}}$  of 4,709 posts from two *troll pages*; in such posts, not rarely satirical jokes are mixed up with serious claims and become debated issues [19]. The most popular post in our collection has 132K shares and says that in the year of the post (2013), after 5,467 years, the month of October has 5 Tuesdays, Wednesdays, and Thursdays, and that this is a very rare event so that Chinese people call it year of the glory *shu tan tzu*. The fact that October 2013



**Fig. 7.** Average polarization of users who liked troll posts (false information) for increasing values of the number of shares. The average polarization increases with the share size, indicating that very popular posts containing false information are mostly supported by conspiracy users.

has 5 Tuesdays, Wednesdays, and Thursdays is true, however the remainder is false: this happens around once every seven years and the phrase has no sense in Chinese. The second most popular post is the popular case of Senator Cirenga mentioned in the introduction, which received more than 36K shares. The third one (shared 28K times) has a more political taste by stating that the former Italian Speaker of the House received a large sum of money after his resignation [52–54].

**Polarization and false information.** In order to understand if polarization facilitates the diffusion of false information, for each post  $p$  in  $\mathcal{P}_{\text{troll}}$  we calculate the number of shares  $x$  of the post and the average polarization of user liking such posts. In Fig. 7 we show the average value of  $\rho$  for increasing levels of shares; more precisely, we compute the average polarization of all the users who liked troll posts with number of shares greater than  $x$ . We find an increasing trend that starts from an average polarization of  $\sim 0.6$  and asymptotically stabilizes at a polarization level  $\sim 0.73$ ; the polarization starts to increase sharply at  $x \sim 20$  and already saturates at  $x \sim 200$ . Figure 7 indicates that false information is more consumed by users with polarization  $\sim 1$  (conspiracy) than users with polarization  $\sim 0$  (science); moreover, the bigger the spreading of the information, the larger the fraction of conspiracy users “liking” it. As an example, the ratio among totally polarized scientific users ( $\rho = 0$ ) and totally polarized conspiracy users ( $\rho = 1$ ) liking information with more than  $2.8 \times 10^4$  shares is  $\sim 1 : 5$ . Hence, users with a conspiracy-like polarization seem to be a more susceptible medium for the diffusion of false information.

## 6 Conclusion

In summary, we find that Facebook users (at least in the Italian dataset) tend to be very polarized with respect to science vs conspiracy subjects, by forming two distinct groups. Such groups are very similar: they present a strong homophily (tendency to

interact with users of similar polarisation), they consume information with similar patterns, and the internal social network structure is statistically similar. However, when it comes to the susceptibility to diffuse false information, we find that users with a polarization toward conspiracy are the most inclined to spread unverified rumors, possibly in a non-linear fashion. In fact, the larger the spread of the false information, the higher the polarization of the supporting users. Therefore, homophily and polarization are possibly the key metrics to target the communities where cascades of false or misleading rumors (the atoms of misinformation) are more likely to spread. Next envisioned steps for this research are a) to consider how unsusceptible users behave with respect to deliberate false information; and b) to analyse whether the topology has a role in content consumption.

Funding for this work was provided by EU FET project MULTIPLEX nr. 317532, EU HOME/2013/CIPS/AG/4000005013 project CI2C, and EU project SIMPOL nr. 610704. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript. Special thanks for building the atlas of pages to *Protesi di Protesi di Complotto, Che vuol dire reale, La menzogna diventa verita e passa alla storia, Simply Humans, Semplicemente me*. Finally, we want to mention Salvatore Previti, Elio Gabalo, Titta, Morocho, and Sandro Forgione for interesting discussions.

## References

1. G. Lotan, E. Graeff, M. Ananny, D. Gaffney, I. Pearce, D. Boyd, *Int. J. Commun.* **5**, 1375 (2011)
2. K. Lewis, M. Gonzalez, J. Kaufman, *Proceedings of the National Academy of Sciences* **109**, 68 (2012)
3. J. Leskovec, D. Huttenlocher, J. Kleinberg, Signed networks in social media, *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*, CHI '10 pp. 1361–1370 (ACM, New York, 2010)
4. J. Kleinberg, *Phil. Trans. R. Soc. A: Mathematical, Physical and Engineering Sciences*, **371**, 20120378 (2013)
5. R. Kahn, D. Kellner, *New Media and Society* **6**, 87 (2004)
6. P.N. Howard, *The Arab Spring's Cascading Effects* (Miller-McCune, 2011)
7. S. Gonzalez-Bailon, J. Borge-Holthoefer, A. Rivero, Y. Moreno, The dynamics of protest recruitment through an online network, *Sci. Rep.* **1**, 197 (2011)
8. R. Bond, J. Fariss, J. Jones, A. Kramer, C. Marlow, J. Settle, J. Fowler, *Nature* **489**, 295 (2012)
9. S. Buckingham Shum, K. Aberer, A. Schmidt, S. Bishop, P. Lukowicz, S. Anderson, Y. Charalabidis, J. Domingue, S. Freitas, I. Dunwell, B. Edmonds, F. Grey, M. Haklay, M. Jelasity, A. Karpıştšenko, J. Kohlhammer, J. Lewis, J. Pitt, R. Sumner, D. Helbing, *Eur. Phys. J. Special Topics* **214**, 109 (2012)
10. D. Centola, *Science* **329**, 1194 (2010)
11. M. Paolucci, T. Eymann, W. Jager, J. Sabater-Mir, R. Conte, S. Marmo, S. Picascia, W. Quattrociocchi, T. Balke, S. Koenig, T. Broekhuizen, D. Trampe, M. Tuk, I. Brito, I. Pinyol, D. Villatoro, *Social Knowledge for e-Governance: Theory and Technology of Reputation*. Roma: ISTC-CNR, 2009
12. W. Quattrociocchi, R. Conte, E. Lodi, *Advances in Complex Systems* **14**, 567 (2011)
13. W. Quattrociocchi, M. Paolucci, R. Conte, *Int. J. Knowledge and Learning* **5**, 457 (2009)
14. V. Bekkers, H. Beunders, A. Edwards, R. Moody, *The Information Society* **27**, 209 (2011)
15. W. Quattrociocchi, G. Caldarelli, A. Scala, Opinion dynamics on interacting networks: media competition and social influence, *Sci. Rep.* **4**, 4938 (2014)
16. L. Howell, Digital wildfires in a hyperconnected world, *WEF Report 2013, World Economic Forum*, 2013

17. Study explains why your stupid facebook friends are so gullible. [http://www.slate.com/blogs/the\\_world\\_/2014/03/18/facebook\\_and\\_misinformation\\_study\\_explains\\_why\\_your\\_stupid\\_facebook\\_friends.html](http://www.slate.com/blogs/the_world_/2014/03/18/facebook_and_misinformation_study_explains_why_your_stupid_facebook_friends.html), March 2014
18. Facebook, Trolls, and Italian Politics. <http://www.globalpolis.org/facebook-trolls-italian-politics/>, March 2014
19. G. Ambrosetti, I forconi: “il senato ha approvato una legge per i parlamentari in crisi”. chi non verrà rieleto, oltre alla buonuscita, si beccherà altri soldi. sarà vero? <http://bit.ly/1JsXXbv>, 8 2013. last checked: 19.01.2014.
20. Ebola Lessons: How Social Media Gets Infected. <http://www.informationweek.com/software/social/-ebola-lessons-how-social-media-gets-infected/a/d-id/1307061>, March 2014
21. The Ebola Conspiracy Theories. <http://www.nytimes.com/2014/10/19/sunday-review/the-ebola-conspiracy-theories.html>, March 2014.
22. The inevitable rise of Ebola conspiracy theories. <http://www.washingtonpost.com/blogs/wonkblog/wp/2014/10/13/the-inevitable-rise-of-ebola-conspiracy-theories/>, March 2014
23. M.L. Meade, H.L. Roediger, *Memory and Cognition* **30**, 995 (2002)
24. C. Mann, F. Stewart, *Internet Communication and Qualitative Research: A Handbook for Researching Online (New Technologies for Social Research series)*. (Sage Publications Ltd, September, 2000)
25. R.K. Garrett, B.E. Weeks, The promise and peril of real-time corrections to political misperceptions, *Proc. 2013 Conf. on Computer supported cooperative work*, Proceedings of CSCW '13 (ACM, 2013), pp. 1047–1058
26. A. Bessi, G. Caldarelli, M. Del Vicario, A. Scala, W. Quattrociocchi, Social determinants of content selection in the age of (mis)information, in *Social Informatics*, Vol. 8851 of Lecture Notes in Computer Science (Springer International Publishing, 2014), pp. 259–268
27. D. Mocanu, L. Rossi, Q. Zhang, M. Karsai, W. Quattrociocchi, Collective attention in the age of (mis)information, in *Computers in Human Behavior* (2015)
28. A. Bessi, M. Coletto, G.A. Davidescu, A. Scala, G. Caldarelli, W. Quattrociocchi, Science vs conspiracy: Collective narratives in the age of misinformation, *PLoS ONE* **10**, 02 (2015)
29. M.D. Vicario, A. Bessi, F. Zollo, F. Petroni, A. Scala, G. Caldarelli, H.E. Stanley, W. Quattrociocchi, *Proceedings of the National Academy of Sciences* **113**, 554 (2016)
30. J. Ugander, B. Lars, M. Cameron, J. Kleinberg, *Proceedings of the National Academy of Sciences* (2012)
31. D. Centola, M. Macy, *Amer. J. Socio.* **113**, 702 (2007)
32. S. Aral, D. Walker, *Science* **337**, 337 (2012)
33. M. McPherson, L. Smith-Lovin, J.M. Cook, *Annu. Rev. Sociol.* **27**, 415 (2001)
34. L.M. Aiello, A. Barrat, R. Schifanella, C. Cattuto, B. Markines, F. Menczer, *ACM Transactions on the Web (TWEB)* **6**, 9 (2012)
35. S. Aral, L. Muchnik, A. Sundararajan, *Proceedings of the National Academy of Sciences* **106**, 21544 (2009)
36. E. Bakshy, I. Rosenn, C. Marlow, L. Adamic, The role of social networks in information diffusion, *WWW'12* (2012)
37. D. Centola, J. Conflict Resolution **51**, 905 (2007)
38. E. Adar, Li Zhang, L.A. Adamic, R.M. Lukose, Implicit structure and the dynamics of blogspace, *Workshop on the weblogging ecosystem*, Vol. 13, pp. 16989–16995 (2004)
39. D. Gruhl, R. Guha, D. Liben-Nowell, A. Tomkins, Information diffusion through blogspace, in *Proc. 13th Int. Conf. on World Wide Web*, pp. 491–501 (ACM, 2004)
40. J. Leskovec, M. McGlohon, C. Faloutsos, N.S. Glance, M. Hurst, Patterns of cascading behavior in large blog graphs, in *SDM*, Vol. 7, pp. 551–556. (SIAM, 2007)
41. B. Golub, M.O. Jackson, *Proceedings of the National Academy of Sciences* **107**, 10833 (2010)
42. D. Liben-Nowell, J. Kleinberg, *Proceedings of the National Academy of Sciences* **105**, 4633 (2008)

43. E. Bakshy, J.M. Hofman, W.A. Mason, D.J. Watts, Everyone's an influencer: quantifying influence on twitter, in *Proc. of the fourth ACM Int. Conf. on Web search and data mining* (ACM, 2011), pp. 65–74
44. K. Lerman, R. Ghosh, Information contagion: An empirical study of the spread of news on digg and twitter social networks. *ICWSM*, Vol. 10, pp. 90–97 (2010)
45. C. Castillo, M. Mendoza, B. Poblete, Information credibility on twitter, in *Proceedings of the 20th international conference on World wide web* (ACM, 2011), pp. 675–684
46. J. Cheng, L. Adamic, P.A. Dow, J.M. Kleinberg, J. Leskovec, Can cascades be predicted? in *Proceedings of the 23rd international conference on World Wide Web*, pp. 925–936. International World Wide Web Conferences Steering Committee, 2014
47. P.A. Dow, L.A. Adamic, A. Friggeri, The anatomy of large facebook cascades, in *ICWSM* (2013)
48. A. Friggeri, L.A. Adamic, D. Eckles, J. Cheng, Rumor cascades, in *Proceedings of the Eighth International AAAI Conference on Weblogs and Social Media*. Association for the Advancement of Artificial Intelligence (2014)
49. Facebook Page: Semplicemente me, November 2014. <https://www.facebook.com/pages/Semplicemente-me/286003974834081?fref=ts>
50. Facebook Page:Simply Humans, November 2014. <https://www.facebook.com/HumanlyHuman?fref=ts>
51. J. Brown, A.J. Broderick, N. Lee, J. Interactive Marketing, **21**, 2 (2007)
52. Troll Post 1. Semplicemente me, November 2014. <https://www.facebook.com/286003974834081/photos/a.286013618166450.51339.286003974834081/403452636422547>
53. Troll Post 2. Semplicemente me, November 2014. <https://www.facebook.com/286003974834081/photos/a.307604209340724.56105.286003974834081/335640103203801>
54. Troll Post 3. Semplicemente me, November 2014. <https://www.facebook.com/286003974834081/photos/a.307604209340724.56105.286003974834081/306244606143351>