Polarization of the Vaccination Debate on Facebook

Ana Lucía Schmidt[a1], Fabiana Zollo[a2], Antonio Scala[b], Cornelia Betsch[c] and Walter Quattrociocchi[a3]

Corresponding author: Ana Lucía Schmidt (analucia.schmidt@unive.it)

[a]Ca' Foscari University of Venice, Via Torino, 155 - 30172, Venice, Italy
Electronic address: [1]analucia.schmidt@unive.it
[2]fabiana.zollo@unive.it
[3]w.quattrociocchi@unive.it
[b]ISC-CNR, SC-CNR Sapienza University of Rome, Via dei Taurini, 19 - 00185, Rome
Electronic address: antonio.scala@phys.uniroma1.it
[c]University of Erfurt, Nordhäuserstr, 63, 9089 Erfurt, Germany
Electronic address: cornelia.betsch@uni-erfurt.de

Abstract

Background

Vaccine hesitancy has been recognized as a major global health threat. Having access to any type of information in social media has been suggested as a potential influence on the growth of anti-vaccination groups. Recent studies w.r.t. other topics than vaccination show that access to a wide amount of content through the Internet without intermediaries resolved into major segregation of the users in polarized groups. Users select information adhering to theirs system of beliefs and tend to ignore dissenting information.

Objectives

The goal was to assess whether users' attitudes are polarized on the topic of vaccination on Facebook and how this polarization develops over time.

Methods

We perform a thorough quantitative analysis by studying the interaction of 2.6M users with 298,018 Facebook posts over a time span of seven years and 5 months. We applied community detection algorithms to automatically detect the emergence of communities accounting for the users' activity on the pages. Also, we quantified the cohesiveness of these communities over time.

Results

Our findings show that the consumption of content about vaccines is dominated by the echo chamber effect and that polarization increased over the years. Well-segregated communities emerge from the users' consumption habits i.e., the majority of users consume information in favor or against vaccines, not both.

Conclusion

The existence of echo chambers may explain why social-media campaigns that provide accurate information have limited reach and be effective only in sub-groups, even fomenting further opinion polarization. The introduction of dissenting information into a sub-group is disregarded and can produce a backfire effect, thus reinforcing the pre-existing opinions within the sub-group. Public health professionals should try to understand the contents of these echo chambers, for example by getting passively involved in such groups. Only then it will be possible to find effective ways of countering anti-vaccination thinking.

Keywords

social media, anti-vaccine sentiment, network analysis, computational social science, misinformation

Introduction

Despite the scientific consensus that vaccines are safe and effective, unsubstantiated claims doubting their safety still occur to this day. Perhaps the most famous case is the multiple times disproved [1,2,3] myth that the MMR vaccine causes autism. However, outbreaks and deaths resulting from objections to vaccines continue to happen [4,5], with the anti-vaccination movement gaining media attention as a result. Mandatory vaccination policies only seem to foment the controversy [6]. Although vaccine hesitancy may have many causes, a lack of confidence is certainly a prominent one [35].

Since 2013, the World Economic Forum has been listing massive digital misinformation among the main threats to our society [7]. Recent studies outline that misinformation spreading is a consequence of the shift of paradigm in content consumption induced by the advent of social media. Indeed, social media platforms like Facebook or Twitter have created a direct path for users to produce and consume content, reshaping the way people get informed [8-13]. Since misinformation influences individuals' beliefs (e.g. risk perceptions), it may also influence the attitude towards vaccination [36]. It has frequently been discussed that social media play a role in the formation of vaccine hesitancy [37].

Like for other misinformation campaigns, Facebook provides an ideal medium for the diffusion of anti-vaccination ideas. Users can access a wide amount of information and narratives and selection criteria are biased toward personal viewpoints [14,15,16]. Online users select information adhering to their system of beliefs, tending to ignore dissenting information and form the so-called echo chambers i.e., polarized groups of like-minded people who keep framing and reinforcing a shared narrative [17,18,19]. The interaction with content dissenting from the shared narrative is mainly ignored and might even foment users segregation, heated debating and, thus, burst opinion polarization [20]. Such a scenario is not limited just to conspiracy theories, but applies to all issues that users perceive as "critical", such as geopolitics or health topics [21] and facilitates the emergence of polarized groups [12] i.e., clusters of users with opposing views that rarely interact with one another.

In this paper, we perform a quantitative analysis to study the evolution of the debate about vaccines on Facebook, taking into account two groups (communities) with opposing views, anti- and pro-vaccine. Considering the liking and commenting behavior of 2.6M users, we analyze the evolution of both communities over time, taking into account the number of users and pages, and their cohesiveness. Our findings confirm the existence of two polarized communities. Additionally, we find evidence that selective exposure plays a pivotal role in how users consume content online. The two communities display different rates of engagement, with the users of the anti-vaccine community being generally more active than those active in the pro-vaccine community.

Methods

The Facebook Platform

Facebook is an online social networking website where people can create profiles or pages to connect with other people and share information such as life events, photos, videos and articles. As of the fourth quarter of 2017, Facebook had 2.2 billion monthly active users. Users on Facebook can interact with posts (i.e., textual content, videos, photos, or links pointing to external documents) from other people or public pages by adding comments or giving a thumbs up (like). More specifically, users' actions allowed by Facebook interaction paradigm are *likes*, *shares*, and *comments*. Each action has a particular meaning [38]: a like represents a positive feedback to a post, a share expresses a desire to increase the visibility of a given information, and a comment is the way in which collective debates take form around the topic of the post.

Ethics Statement

The data collection process was carried out using the Facebook Graph API [22], which is publicly available. The pages from which we downloaded data are public Facebook entities and can be accessed by anyone. Users' content contributing to such pages is also public unless users' privacy settings specify otherwise, and in that case it is not available to us.

Data Collection

The dataset was generated using the Facebook Graph API to search for pages containing the keywords *vaccine*, *vaccines* or *vaccination* in their name or description. We then cleaned the raw Facebook results. Inclusion criteria were language (English), a minimum level of activity on the page (at least 10 posts), date of the posts (between 1st January 2010 to 31st May 2017), and relation of the page to the topic of vaccination. This last step was essential, since having one of the keywords in the description does not necessarily mean the page's topic is about vaccines. False positive search results are, for example, the pages *The Vaccines* (an UK music band) or *Arthur D'vaccine* (a comedian).

From the resulting set of Facebook pages, we used the Graph API to download all the posts as well as all the related likes[1] and comments. Considering the narrative of the pages and the content of the posts, all the Facebook pages were also manually classified by two raters into two main groups: 145 *pro-vaccine* with 1,388,677 users and 98 *anti-vaccine with* 1,277,170 users. The Cohen's kappa inter-agreement between both raters is 0.966, showing nearly perfect agreement. All the authors approved and verified the final classification. The complete list of the Facebook pages with their respective community label and a breakdown of the dataset are reported in the Appendix.

Preliminaries and Definitions

In network theory a *bipartite network* is a graph whose vertices can be divided into two disjoint and independent sets. The likes (or comments) given by users to the posts of different Facebook pages form a *bipartite network*. This *bipartite network* is formed by a set of users and a set of pages where links only exist between a user and a page if the user liked (or commented) anything on that page.

We can represent the bipartite network with a matrix where each column is a user and each row is a page, and the content of each cell equals 1 if the user liked a post of that page, and 0 otherwise. If we multiply the matrix

[1] Since Facebook started introducing reactions (love, haha, wow, sad, angry) in February 2016, only the likes were considered for the whole period.

by its transpose, we get the *projection of the bipartite network*. This new matrix will have a row and column for each page, and the content of each cell will represent the number of common users between the 2 pages that define that cell, that is, the number of users who liked any post on both pages. The same method can also be applied considering the matrix formed by the users' comments.

For illustration, Figure 1 visualizes a simplified example of a bipartite network with 5 users and 4 pages and the corresponding projection.

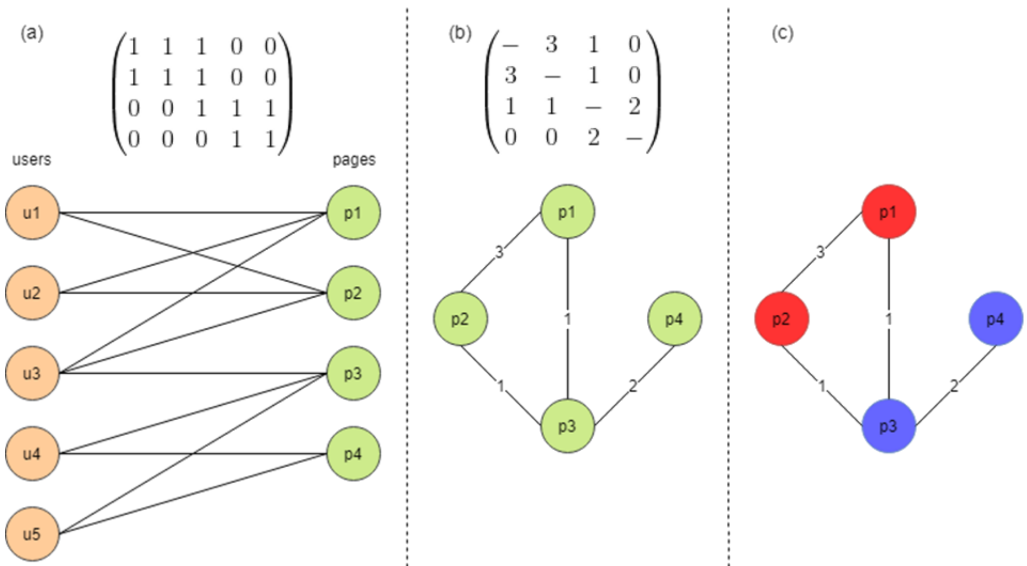

Figure 1 - (a) Bipartite network with 5 users and 4 pages, the links between them indicate that a user liked a page. (b) The projection of the bipartite network, the weights on the links between the pages show the number of users they have in common. (c) The community structure as detected with the algorithm FastGreedy. When nodes share a color they belong to the same community.

Once we have the network of pages linked by their common users (Figure 1b), we can apply different community detection algorithms to detect *communities*, groups of pages that are strongly connected (Figure 1c). To do this we apply five well known community detection algorithms: FastGreedy[2][23], WalkTrap[3][24], MultiLevel[4][25] and

[2] It optimizes the modularity score in a greedy manner to calculate the communities. The modularity is a benefit function that measures the quality of a particular division of a network into communities. A high modularity score corresponds to a dense connectivity between nodes inside a cluster and sparse connections between clusters. This algorithm takes an agglomerative bottom-up approach: initially each vertex belongs to a separate community and, at each iteration, the communities are merged in a way that yields the largest increase in the current value of modularity.

[3] It exploits the fact that a random walker tends to become trapped in the denser parts of a graph i.e, in communities.

[4] It uses a multi-level optimization procedure for the modularity score. It takes a bottom-up approach where each vertex initially belongs to a separate community and in each step, unlike FastGreedy, vertices are reassigned to a new community.

LabelPropagation[5][26]. Different algorithms are used as they allow for unsupervised clustering i.e., no human intervention, and they each have different approaches to detecting of communities in the networks. To compare the communities detected with the various algorithms we use standard methods that compute the similarity between different community partitions by considering how the algorithms assign the nodes to each community [27]. Due to its speed and its lack of parameters in need of tuning, the FastGreedy algorithm will be the main reference to compare against the partitions resulting from the application of other community detection algorithms. Starting from the communities that emerge from users' behavior, in the following sections our aim is to measure i) the number of pages users from each community interactive with (*selective exposure*), ii) the activity of the users across the communities (*polarization*), and iii) the growth of the communities over time.

Results and Discussion

Validation of the Community Partition

In order to validate the manual partition of the pages into two communities we generated the projections of the bipartite networks considering the user likes and the user comments. We then applied the community detection algorithms to extract the communities of pages according to the users' behavior and compared those to the manual partition.

Table 2 shows the comparison between a random partition of the pages, the manual partition, and the FastGreedy partition against those resulting from the different algorithms. We can see that the manual classification matches well against the unsupervised approaches and that the FastGreedy results have a high agreement with the other algorithms. This indicates that the users' behavior generates well defined communities of pages and that these communities are similar to the anti-vaccine and pro-vaccine communities as manually tagged.

[5] It uses a simple approach where each vertex is assigned a unique label, which is updated according to majority voting in the neighboring vertices. Dense node groups quickly reach a consensus on a common label.

Table 1 – Validation of the community partition.

Graph	Communities	FastGreedy	WalkTrap	MultiLevel	LabelProp.
Likes	**Random**	0.496	0.497	0.495	0.497
	Manual	0.774	0.721	0.738	0.714
	FastGreedy	1	0.935	0.950	0.901
Comments	**Random**	0.497	0.499	0.495	0.496
	Manual	0.590	0.610	0.567	0.570
	FastGreedy	1	0.909	0.876	0.824

Note: We compared a random partition of the pages into communities, the manual classification, and the FastGreedy classification against the community partitions detected with the different community detection algorithms. The values of the comparison range from 0 to 1, where 1 is an exact match and 0 is no match.

Thus, the pages cluster together according to the users' activity. In a next step, we analyzed the polarization of the users.

Polarization

Assuming that a user u has performed x likes on community C1 and y likes on community C2, we calculate the users' polarization as $\rho(u) = (x - y)/(x + y)$. Thus, a user u for whom $\rho(u) = -1$ is polarized towards C2, whereas a user whose $\rho(u) = 1$ is polarized towards C1. We then measure the polarization of all users considering the communities they commented and liked content on. We examine two partitions: the manual classification of pages, pro-vaccine and anti-vaccine, and the two biggest communities as detected with FastGreedy, C1 and C2.

Figure 2 shows the Probability Density Function (PDF) of $\rho(u)$ for all users who have given at least 10 likes in their lifetime. The PDF for the polarization of all users is sharply bi-modal, that is, the majority of the users are either at -1 or at 1. This indicates a strong polarization among the communities, that is, the majority of the users are active either in the pro-vaccine or anti-vaccine community, not both.

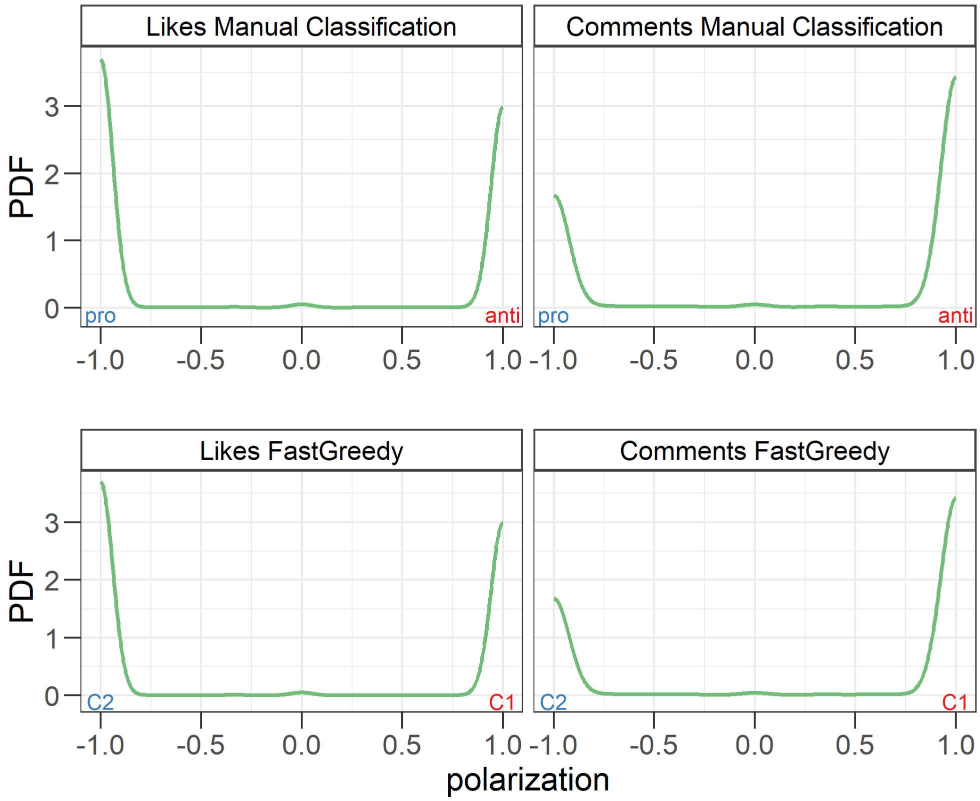

Figure 2 - Probability Density Function (PDF) of the users' liking (left) and commenting (right) behavior in the manual communities (top) and the 2 largest communities detected with FastGreedy (bottom). The distribution of the users is bimodal for all cases, which indicates a strong polarization among the communities, that is, the majority of the users are active in only one community.

Selective Exposure

Facebook users differ in the time they spend with the pages and in how frequently they interact with the pages. The *lifetime* of a user is defined as the period of time where the user started and stopped interacting with the included set of pages. It can be approximated by the time difference between a user's latest and earliest liked post. The total number of likes per user is a good proxy for the user's *activity* i.e., their level of engagement with the Facebook news pages. These two measures can provide important insights on how users consume information in each echo chamber as demonstrated in the following analyses.

Figure 3 shows the number of unique pages users from the anti-vaccine (red) and pro-vaccine communities (blue) interact with, considering increasing levels of lifetime and activity for different time windows (yearly left, monthly middle and weekly right panel). For a comparative analysis, we standardized lifetime and activity to range between 0 and 1, both over the entire user set of each community, and the number of pages.

Note that for both communities, users usually interact with a small number of Facebook pages. Longer lifetime and higher levels of activity correspond with less number of pages being consumed. This suggests that more time on Facebook corresponds to a smaller variety of sources being consumed. This is consistent with [12] showing that content consumption on Facebook is dominated by selective exposure and, over time, users personalize their information sources accordingly with their tastes which results in a smaller number of sources being consumed.

Pro-vaccine users interact with M = 1.42 pages (SD = 0.79), anti-vaccine users with 2.45 (SD = 2.13). This difference is displayed in Figure 3: users in the anti-vaccine community (red line) consume information from a more diverse set of pages than those in the pro-vaccine community, regardless of the time window considered. Grey shades are 95% CI of the local regression of the data, indicating significant differences between the groups at any time. So while there is a natural tendency of users to confine their activity to a limited set of pages [12], the two communities display different rates of selective exposure. The anti-vaccine community shows more commitment to the consumption of their posts.

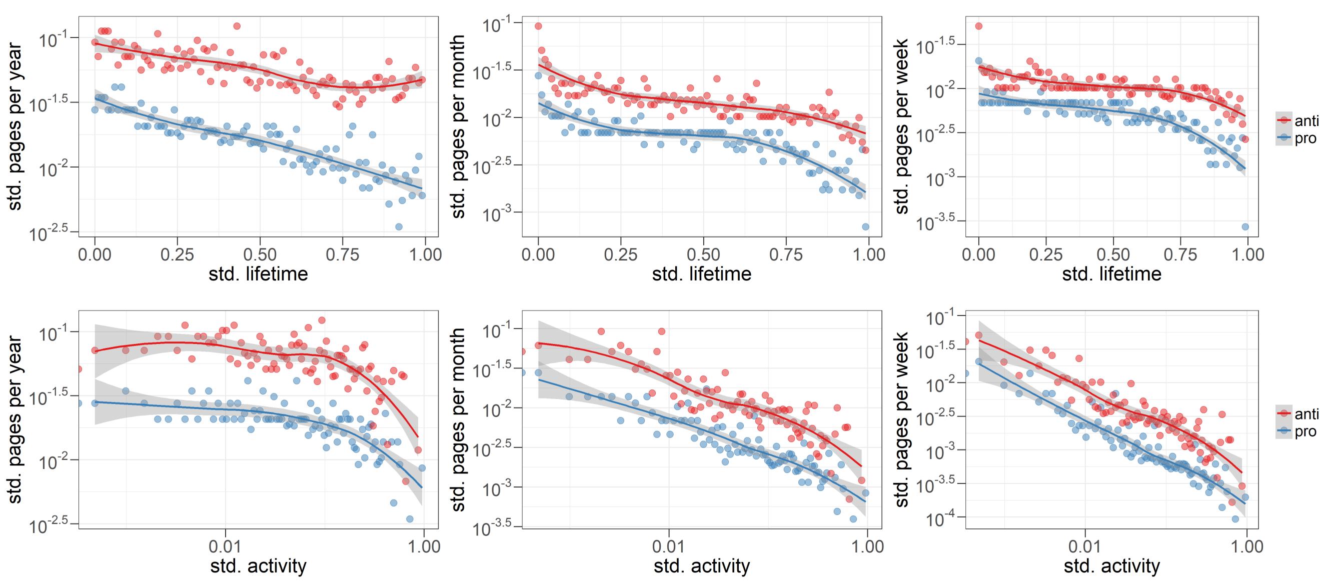

Figure 3 - Maximum number of unique pages that users with increasing levels of standardized lifetime (top) or standardized activity (bottom) interact with yearly (left), monthly (middle) and weekly (right) for each community. Users' lifetime corresponds to the standardized time difference between their latest and earliest liked post. Users' activity corresponds to the standardized number of likes given in their lifetime. Users display a tendency to like less pages when their lifetime and activity increases. The users who interact with the anti-vaccine community also consume a larger variety of pages than the pro-vaccine users. Grey shades are 95% CI of the fitted curve, indicating significant differences between the groups at any time.

Growth of the Communities over Time

We also analyzed the growth of the two communities over time, considering the variety of pages and the number of users that interact with them. Figures 4 shows the evolution of the communities over the years in quarterly increments.

The left panel plots the number of active pages in each community. We define a page as active in a specific quarter if it made a post (bottom panel), received a like (middle) or comment in that period (upper panel). Overall, the number of active pages in both communities increases at similar rates, with slight variations when we consider the different types of action that marks a page as active. If we use the pages' posting activity or the likes they received to determine whether they were active in a given quarter, we can see that, from 2013, the pro-vaccine community consistently tends to show a higher number of active pages than the anti-vaccine community (interaction effect in a MANOVA with sentiment (pro, anti) and time (until 2012Q4 vs. following) as factors and posts and likes as dependent variables $F(2,55) = 2.708$, $p = 0.076$; $eta^2 = 0.09$; both main effects are highly significant). On the other hand, if we focus on the comments, the anti-vaccine community shows a boost in activity starting in 2015 (interaction effect in an ANOVA with sentiment (pro, anti) and time (until 2014Q4 vs. following) as factors and comments as dependent variable $F(1,56) = 5.053$, $p = 0.029$; $eta^2 = 0.08$; both main effects are significant).

The right panel plots the number of active users in each community. We define users as active if they gave a like (or comment) to any page of that community in the given quarter. The plot shows that while both communities gain users throughout the entire period, the anti-vaccine community has, until the end of 2015 and beginning of 2016, more active users than the pro-vaccine community. After that, this relation reverses (interaction effect in a MANOVA with sentiment (pro, anti) and time (until 2015Q4 vs. following) as factors and comments and likes as dependent variables $F(2,55) = 12.218$, $p < 0.001$; $eta^2 = 0.31$; both main effects are highly significant).

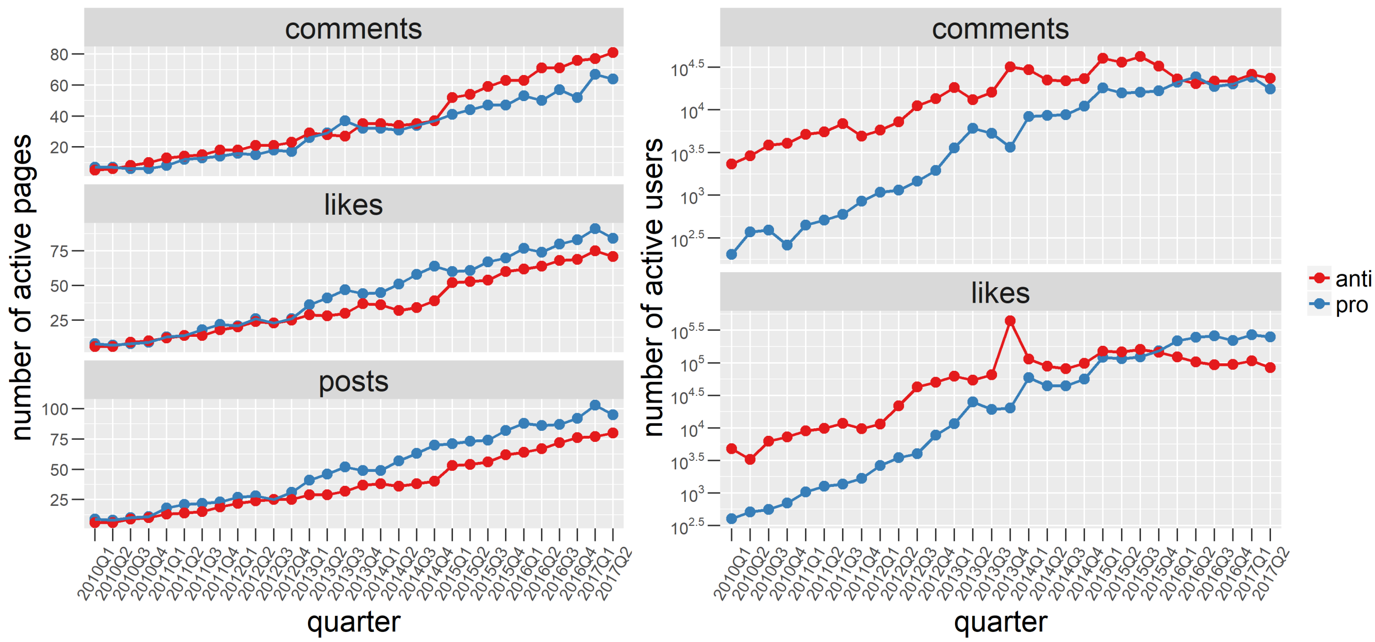

Figure 4 – Number of active pages (left) and users (right) in each community. We define a page as active in a specific quarter if it made a post (bottom panel), received a like (middle panel) or comment (upper panel) in that period. We define a user as active in a community on a given quarter, if they gave a like (bottom panel) or comment (top panel) to any page of that community in that time.

Another important factor to consider is the cohesiveness of the pro-vaccine and anti-vaccine communities. In order to analyze whether the growth of the communities depends on the emergence of isolated pages or whether it grows steadily, we split the projections of the bipartite likes and comments graph by the community of the pages. This results in 4 sub-graphs, each containing the pages of one community, pro-vaccine or anti-vaccine, and the common users that linked them considering the likes or the comments. We can then calculate the fragmentation of each community by applying the community detection algorithms and obtaining their partition.

Figure 5 shows the number of pages of the biggest sub-community of the anti-vaccine (left) or pro-vaccine communities (right) in a given quarter, that is, the largest connected component found with the different community detection algorithms. The black line represents the total number of pages in the sub-graphs in that quarter. It marks the maximum possible size for the largest connected component to take in that moment in time. The closer the size of the largest connected component is to the total number of pages in the system, the more tightly linked that community is in that moment in time.

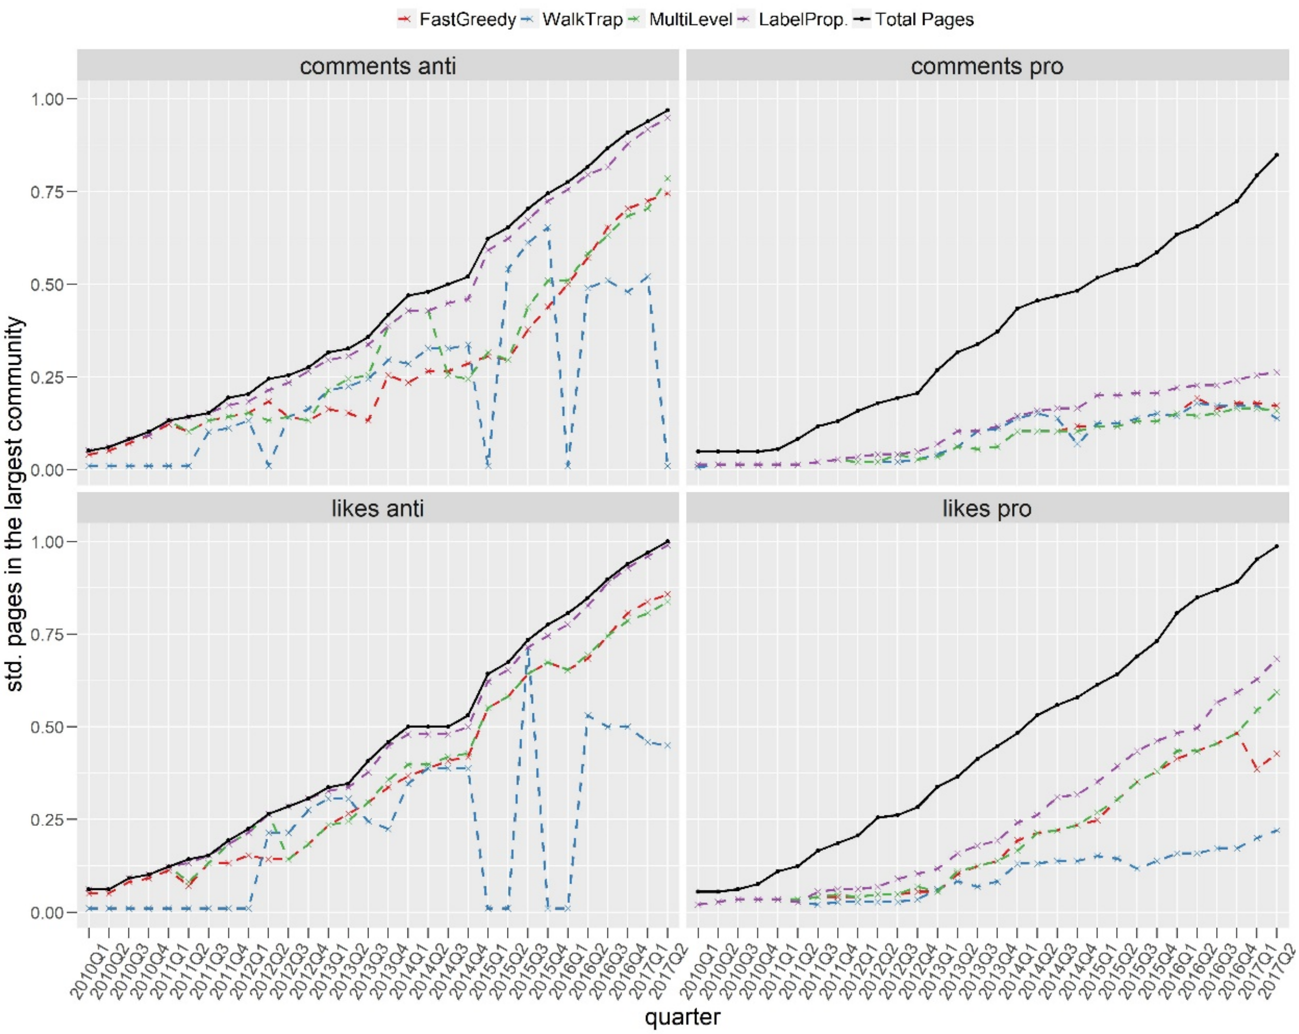

Figure 5 - Size of the largest connected component within the set of pages tagged as anti-vaccine and pro-vaccine over time, considering various community detection algorithms. The black line represents the total number of pages over time in the anti-vaccine and pro-vaccine communities, that is, the maximum possible size for the largest connected component in that moment in time. The graph shows that the anti-vaccine community grows cohesively, with the new pages joining the already existing group of pages, while the pro-vaccine community grows in a more fragmented, independent way.

The plots show that in the anti-vaccine community the number of pages in the largest component remains close to the total number of pages in the system. In the case of the pro-vaccine sub-graphs, however, the size of the largest community does not increase closely with the number of pages in the system. This signifies that the anti-vaccine community grows in a more cohesive manner, with pages tightly linked by their users' activity, while the pro-vaccine community is more fragmented.

Discussion

By means of quantitative analysis of Facebook likes and comments we validated the existence of two opposing narratives regarding the vaccination debate on Facebook. We show that the communities emerge from the users' consumption habits and that users are highly polarized, that is, the majority of users only consumes and produces information in favor or against vaccines, not both.

We also showed that both narratives are subjected to selective exposure, and that the more active a user is on Facebook the smaller is the variety of sources they tend to consume. We note, however, that the users from the anti-vaccination community consume more sources compared to the pro-vaccine users. This is consistent with the results of previous studies [14] that show that people in conspiracy-like groups show higher engagement with the community. One can (very carefully) conclude that anti-vaccination attitudes are rooted more deeply in the social and psychological background of a person than pro-vaccination.

We also analyzed the communities' evolution over time. While the pro-vaccine pages are generally more active, the anti-vaccine pages concentrate the majority of the debate, receiving more comments from users. We show that the anti-vaccine community had a more active user base until the end of 2015, where the activity seems to stall. This matches with the outbreak of measles at Disneyland [4], which put the anti-vaccination movement in the spotlight and gained the attention of mainstream media [28-34]. Further studies are needed to determine the reason for this stagnation.

Finally, we show that while both narratives have gained attention on Facebook over time, anti-vaccine pages display a more cohesive growth (i.e. pages are liked by the same people), while the pro-vaccine pages seem to grow in a highly fragmented fashion (i.e. pages are liked by different people).

Limitations

The data collection process was done the 5th of June 2017 and represents a snapshot of the pages, posts, comments and likes available at the time. Pages, posts, likes and comments that were made in the downloaded period (1st January 2010 to 31st May 2017) and were removed before the download date are not present in the dataset. The data only includes the likes and comments by users whose privacy settings allowed public access to their activity on public pages on the download date.

Conclusions

Facebook allows echo chambers to emerge, in which pro- and anti-vaccination attitudes polarize the users. Social media campaigns that use Facebook to advocate for vaccination and provide accurate information should expect to only reach pro-vaccination users as there is nearly no interaction between the groups. Overall, social media seem to be a powerful promoter of different sentiments about vaccination and therefore it is likely that it contributes to vaccine avoidance.

Appendix

Table 2 - Dataset Description.

	Anti-vaccine	Pro-vaccine
Pages	98	145
Posts	189,759	108,259
Likes	12,696,440	11,459,295
Likers	1,145,650	1,325,511
Comments	1,351,839	749,209
Commenters	271,598	146,196
Users	1,277,170	1,388,677

Note: The posts, likes and comments are considered pro or anti vaccines if they were made on a page classified as such. Likers is the number of unique users who have given at least one like to the community. Commenters is the unique number of users who have given at least one comment to the community. Users is the number people who have given at least a like or a comment to the community.

References

[1] Chen, W., Landau, S., Sham, P., & Fombonne, E. (2004). No evidence for links between autism, MMR and measles virus. Psychological medicine, 34(3), 543-553.

[2] DeStefano, F. (2007). Vaccines and autism: evidence does not support a causal association. Clinical Pharmacology & Therapeutics, 82(6), 756-759.

[3] American Academy of Pediatrics (2013) Vaccine safety: Examine the evidence. American Academy of Pediatrics.

[4] Zipprich, J., Winter, K., Hacker, J., Xia, D., Watt, J., Harriman, K., & Centers for Disease Control and Prevention (CDC). (2015). Measles outbreak—California, December 2014-February 2015. MMWR Morb Mortal Wkly Rep, 64(6), 153-154.

[5] Bocci, E., & Naselli, M. (2017). Morbillo, muore una bambina di 9 anni a roma. La Repubblica, June 2017.

[6] Betsch, C., & Böhm, R. (2015). Detrimental effects of introducing partial compulsory vaccination: experimental evidence. The European Journal of Public Health, 26(3), 378-381.

[7] Quattrociocchi, W. (2017). Part 2-Social and Political Challenges: 2.1 Western Democracy in Crisis?. In World Economic Forum.

[8] Brown, J., Broderick, A. J., & Lee, N. (2007). Word of mouth communication within online communities: Conceptualizing the online social network. Journal of interactive marketing, 21(3), 2-20.

[9] Kahn, R., & Kellner, D. (2004). New media and internet activism: from the 'Battle of Seattle' to blogging. New media & society, 6(1), 87-95.

[10] Quattrociocchi, W., Caldarelli, G., & Scala, A. (2014). Opinion dynamics on interacting networks: media competition and social influence. Scientific reports, 4, 4938.

[11] Kumar, R., Mahdian, M., & McGlohon, M. (2010, July). Dynamics of conversations. In Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining (pp. 553-562). ACM.

[12] Schmidt, A. L., et al. (2017). Anatomy of news consumption on Facebook. Proceedings of the National Academy of Sciences, 201617052.

[13] Betsch, C., Brewer, N.T., Brocard, P., Davies, P., Gaissmaier, W., Haase, N., Leask, J., Renkewitz, F., Renner, B., Reyna, V.F., Rossmann, C., Sachse, K., Schachinger, A. & Siegrist, M. (2012). Opportunities and Challenges of Web 2.0 for Vaccination Decisions. Vaccine, 30, 3727-3733.

[14] Bessi, A., Coletto, M., Davidescu, G. A., Scala, A., Caldarelli, G., & Quattrociocchi, W. (2015). Science vs conspiracy: Collective narratives in the age of misinformation. PloS one, 10(2), e0118093.

[15] Bessi, A., et al. (2016). Homophily and polarization in the age of misinformation. The European Physical Journal Special Topics, 225(10), 2047-2059.

[16] Mocanu, D., Rossi, L., Zhang, Q., Karsai, M., & Quattrociocchi, W. (2015). Collective attention in the age of (mis) information. Computers in Human Behavior, 51, 1198-1204.

[17] Del Vicario, M., et al. (2016). The spreading of misinformation online. Proceedings of the National Academy of Sciences, 113(3), 554-559. http://www.pnas.org/content/113/3/554.full

[18] Quattrociocchi, W., Scala, A., & Sunstein, C. R. (2016). Echo chambers on Facebook. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2795110

[19] Del Vicario, M., Vivaldo, G., Bessi, A., Zollo, F., Scala, A., Caldarelli, G., & Quattrociocchi, W. (2016). Echo Chambers: Emotional Contagion and Group Polarization on Facebook. Scientific Reports, 6. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5131349/

[20] Zollo, F., et al. (2017). Debunking in a world of tribes. PloS one, 12(7), e0181821.

[21] Bessi, A., Zollo, F., Del Vicario, M., Scala, A., Caldarelli, G., & Quattrociocchi, W. (2015). Trend of Narratives in the Age of Misinformation. PloS one, 10(8), e0134641. http://journals.plos.org/plosone/article?id=10.1371/journal.pone.0134641

[22] Facebook (August 2013) Using the graph API (Facebook, Menlo Park, CA). Available at https://developers.facebook.com/docs/graph-api/using-graph-api. Accessed January 19, 2014.

[23] Clauset, A., Newman, M. E., & Moore, C. (2004). Finding community structure in very large networks. Physical review E, 70(6), 066111.

[24] Pons P, Latapy M (2006) Computing communities in large networks using random walks. J Graph Algorithms Appl 10(2):191–218.

[25] Blondel, V. D., Guillaume, J. L., Lambiotte, R., & Lefebvre, E. (2008). Fast unfolding of communities in large networks. Journal of statistical mechanics: theory and experiment, 2008(10), P10008.

[26] Raghavan, U. N., Albert, R., & Kumara, S. (2007). Near linear time algorithm to detect community structures in large-scale networks. Physical review E, 76(3), 036106.

[27] Rand, W. M. (1971). Objective criteria for the evaluation of clustering methods. Journal of the American Statistical association, 66(336), 846-850.

[28] Salzberg, S. (2015). Anti-vaccine movement causes worst measles epidemic in 20 years. Forbes, February 2015.

[29] Pease, B. (2015). Fear vs. fact: The modern anti-vaccination movement. Harvard Science Review, March 2015.

[30] Haberman, C. (2015). A discredited vaccine study’s continuing impact on public health. The New York Times, February 2015.

[31] Healy, J., &Paulson, M. (2015). Vaccine critics turn defensive over measles. The New York Times, January 2015.

[32] Kristof, N. (2015). The dangers of vaccine denial. The New York Times, February 2015.

[33] Barbash, F. (2015) Disneyland measles outbreak strikes in antivaccination hotbed of California. The Washington Post, January 2015.

[34] Gumbel, A. (2015). US states face fierce protests from anti-vaccine activists. The Guardian, April 2015.

[35] Schmid, P., Rauber, D., Betsch, C., Lidolt, G. and Denker, M.L., 2017. Barriers of influenza vaccination intention and behavior–a systematic review of influenza vaccine hesitancy, 2005–2016. PloS one, 12(1), p.e0170550.

[36] Betsch, C., Böhm, R., & Chapman, G. B. (2015). Using behavioral insights to increase vaccination policy effectiveness. Policy Insights from the Behavioral and Brain Sciences, 2(1), 61-73.

[37] Betsch, C., Brewer, N.T., Brocard, P., Davies, P., Gaissmaier, W., Haase, N., Leask, J., Renkewitz, F., Renner, B., Reyna, V.F. and Rossmann, C., 2012. Opportunities and challenges of Web 2.0 for vaccination decisions. Vaccine, 30(25), pp.3727-3733.

[38] Ellison NB, Steinfield C, Lampe C. The Benefits of Facebook "Friends:" Social Capital and College Stu- dents' Use of Online Social Network Sites. Journal of Computer-Mediated Communication. 2007; 12(4):1143–1168.