

Mapping Social Dynamics on Facebook: The Brexit Debate

Michela Del Vicario^{1,*}

CSSLab, IMT School for Advanced Studies, Lucca, Italy

Fabiana Zollo^{1,*}

*Dept. of Environmental Sciences, Informatics and Statistics
Ca' Foscari University of Venice, Venice, Italy*

Antonio Scala*

ICS CNR, Rome, Italy

Guido Caldarelli*

CSSLab, IMT School for Advanced Studies, Lucca, Italy

Walter Quattrociocchi*

CSSLab, IMT School for Advanced Studies, Lucca, Italy

Abstract

Nowadays users get informed and shape their opinion through social media. However, the disintermediated access to contents does not guarantee quality of information. Selective exposure and confirmation bias, indeed, have been shown to play a pivotal role in content consumption and information spreading. Users tend to select information adhering (and reinforcing) their worldview and to ignore dissenting information. This pattern elicits the formation of polarized groups – i.e., echo chambers – where the interaction with like-minded people might even reinforce polarization. In this work we address news consumption around Brexit in UK on Facebook. In particular, we perform a massive analysis on more than 1 Million users interacting with Brexit related posts from

*Corresponding author

Email address: `michela.delvicario@imtlucca.it` (Michela Del Vicario)

¹These authors contributed equally to this work.

the main news providers between January and July 2016. We show that consumption patterns elicit the emergence of two distinct communities of news outlets. Furthermore, to better characterize inner group dynamics, we introduce a new technique which combines automatic topic extraction and sentiment analysis. We compare how the same topics are presented on posts and the related emotional response on comments finding significant differences in both echo chambers and that polarization influences the perception of topics. Our results provide important insights about the determinants of polarization and evolution of core narratives on online debating.

Keywords: Collective Debates, Polarization, Online Social Networks

1. Introduction

The Arab Spring and Ukrainian revolution showed social media as a liberalizing technology and powerful vehicle of information, engagement, mobilization, able to encourage innovation and democracy. But social media have also
5 changed the way we get informed and form our opinions.

According to a recent report [1], approximately 63% of users acquire their news from social media, and these news stories undergo the same popularity dynamics as other forms of online contents (such as selfies and cat photos). As a result of disintermediated access to information and of algorithms used in
10 content promotion, communication has become increasingly personalized, both in the way messages are framed and how they are shared across social networks.

Selective exposure and confirmation bias, indeed, have been shown to play a pivotal role in content consumption and information spreading [2]. Users tend to select information adhering (and reinforcing) their worldview and to ignore
15 dissenting information [3, 4, 5, 6]. This pattern elicits the formation of polarized groups – i.e., echo chamber – where the interaction within like-minded people might even reinforce polarization [7, 8].

Several studies pointed out the effects of social influence online [9, 10, 11, 12]. Results reported in [13] indicate that emotions expressed by others on Facebook

20 influence our own emotions, providing experimental evidence of massive-scale
contagion via social networks. Recent works [6, 5], indeed, showed that attempts
to debunk false information are largely ineffective. In particular, the discussion
degenerates when the two polarized communities interact with one another.
The increasing interest in online debates led researchers to investigate many
25 of their aspects, from the characterization of conversation threads [14] to the
detection of bursty topics on microblogging platforms [15], to the disclosure of
the mechanisms behind information diffusion for different kinds of contents [16].

More recently, several doubts about social influence on the Internet have
been raised during Brexit –the British referendum to leave the European Union–
30 campaign, where both sides, Leave and Remain, battled it out on social media.
Indeed, a big effort has been dedicated to characterize the dynamics of the
online Brexit debate, from applying filtering algorithms to study the shape of
online data [17], through the investigation of the role of *bots* on the direction of
discussions [18], to the study of the effects of the referendum result on financial
35 markets [19].

In this paper we address the Brexit discussion on Facebook public pages
referring to UK based official information sources listed in the European Media
Monitor [20].

Firstly, we characterize the structural properties of the discussion by ob-
40 serving the spontaneous emergence of two well-separated communities; indeed,
connections among pages are the direct result of users’ activity, and we do not
perform any categorization of contents a priori. Then, we explore the dynamics
behind discussion: looking at users polarization towards the two communities
and at their attention patterns, we find a sharply bimodal distribution, showing
45 that users are divided into two main distinct groups and confine their attention
on specific pages.

Finally, to better characterize inner group dynamics, we introduce a new
technique which combines automatic topic extraction and sentiment analysis.
We compare how the same topics are presented on posts and the related com-
50 ments, finding significant differences in both echo chamber and that polarization

on the perception of topics. We first measure the distance between how a certain concept is presented on the posts and then the emotional response of users to such controversial topics. Our new metrics could be of great interest to identify the most crucial topics in online debates. Indeed, it is highly likely that the
55 greater the emotional distance between the same concept in two echo chambers, the greater the polarization of users involved in the discussion. Therefore, such a distance may be a key marker to locate controversial topics and to understand the evolution of the core narratives within distinct echo chambers.

2. Methods

60 *Ethics Statement.*

The data collection process was carried out using the Facebook Graph API [21], 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
65 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 European Media Monitor (EMM) [20] provides a list of all news sources
70 which includes, for each of them, the related country and region. We limited our collection to all pages whose legal head office (at least one of them) is located in the United Kingdom. For each page, we downloaded all the posts from January 1st to July 15th, 2016, as well as all the related likes and comments. The exact breakdown of data is provided in Tab. 2, while the complete set of downloaded
75 pages is reported in Tab. 5 in *Appendix B*.

Preliminaries and Definitions.

In this section we provide a brief description of the main concepts and tools used in the analysis.

	Total	Brexit
<i>Pages</i>	81	38
<i>Posts</i>	303,428	5,039
<i>Likes</i>	186,947,027	2,504,956
<i>Comments</i>	38,182,541	469,397
<i>Likers</i>	30,932,388	1,365,821
<i>Commenters</i>	7,222,273	259,078

Table 1: **Dataset description:** from January 1st to July 15th, 2016.

Bipartite Projection.

80 A bipartite graph is a triple $\mathcal{G} = (A, B, E)$ where $A = \{a_i \mid i = 1 \dots n_A\}$ and $B = \{b_j \mid j = 1 \dots n_B\}$ are two disjoint sets of vertices, and $E \subseteq A \times B$ is the set of edges – i.e. edges exist only between vertices of the two different sets A and B . The bipartite graph \mathcal{G} is described by the rectangular matrix M defined as

$$M_{ij} = \begin{cases} 1 & \text{if an edge exists between } a_i \text{ and } b_j \\ 0 & \text{otherwise} \end{cases}.$$

85 We consider the bipartite network $\mathcal{G} = (P, U, E)$ where P is the set of Facebook pages concerned on Brexit topics (see Tab. 5 in *Appendix B*) and U is the set of users active on pages belonging to P . An interaction with a given information posted by a page $p \in P$ determines a link between a user $u \in U$ and the page p , hence $M_{p,u} = 1$ indicates that user u was active on page p . For our analysis we use the co-occurrence matrices $C^P = MM^T$ and $C^U = M^T M$ that count, respectively, the number of common neighbors between two vertices of P or U . As an example, $C_{p,q}^P$ for $p \neq q$ counts the number of users that were active on both pages p and q . C^P can be interpreted as the weighted adjacency matrix of the co-occurrence graph G^P with vertices on P . Each non-zero element $C_{p,q}^P$ 95 corresponds to an edge (p, q) among vertices p and q with weight C_{pq}^P .

Community Detection Algorithms.

Community detection algorithms serve to identify groups of nodes in a network. Most of the strategies relies on the modularity which quantifies the division of a network in separated clusters, high modularity corresponds to a dense connectivity between nodes in a community and sparse connections between modules. In this work we apply four different community detection algorithms: Fast Greedy (FG), that seeks for the maximum modularity score by considering all possible community structures in the network. It tries to optimize the modularity function in a greedy manner [22]. Walktrap (WT), that exploits the fact that a random walker tends to remain trapped in the denser part – i.e., communities – of a graph. Hence WT uses short random walks to merge separate communities [23]. Multilevel (ML), that is based on a multi-level modularity optimization procedure [24]. Label Propagation (LP) [25], that is a nearly linear time algorithm that gives unique labels to vertices that are then updated according to majority voting in the neighboring vertices. Dense group of nodes reach consensus on a common label quickly. To compare the various community structures we use standard methods that compute the similarity between different clustering methods by considering how nodes are assigned by each community detection algorithm [26, 27].

Backbone Detection Algorithm.

The disparity filter algorithm is a network reduction technique based on the local identification of the statistically relevant weight heterogeneities. This method is able to identify the backbone structure of a weighted network without destroying its multi-scale nature [28]. We make use of this algorithm to obtain the relevant connections that form our networks' backbones and produce clearer visualizations.

Results and Discussion

As a preliminary step, we divide all UK based pages in two groups: *Brexit pages*, that includes those pages engaged in the debate around the Brexit, and

125 *Non Brexit pages.* Out of 81 pages, 38 posted at least one news story about
the Brexit. Hence, we characterize the users behavior on Brexit pages and their
related posts.

2.1. Communities and News Polarization

Online social media proved to facilitate the aggregation of individuals in
130 communities of interest, also known as *echo chambers* [3], especially when re-
stricting the interaction of users to conflicting information [2, 5]. In the case
of Brexit pages we focus of the emerging communities without considering the
shared contents, but rather by accounting for the connections created by users
activities.

135 Therefore, we start by analyzing the community structure of the Brexit pages
graph. We consider the bipartite projection of the pages-users graph G_p where
nodes are Brexit pages and two pages are connected if at least one user liked a
post from each of them. The weight of a link is determined by the number of
users in common between the two pages.

140 In Fig. 1(a) we show the backbone structure of G_p . Colors (resp., blue and
red) represent the membership to one of the two communities (resp., C1 and
C2) detected by the Fast Greedy (FG) algorithm (see *Methods* section for fur-
ther details). Fig. 1(b) reports the percentage of pages in both communities. A
complete list of the pages and the relative membership is reported in Tab. 5 in
145 *Appendix B*. We compare the results of FG and two other community detection
algorithms i.e., Walktrap (WT) and Multilevel (ML) (refer to *Methods* section
for further details) by means of the Rand method [26, 27]; we find a very high
concordance between FG and ML (0.90), and lower ones between FG and WT
(0.69), and between ML and WT (0.63)². Our analysis underlines the sponta-
150 neous emergence of two separate communities active on Brexit pages, where the
connections among pages are a simple result of the interaction of users on them.

²For further details on the comparison between the community detection algorithms, and
for a complete analysis of the WT case, refer to *Appendix A*.

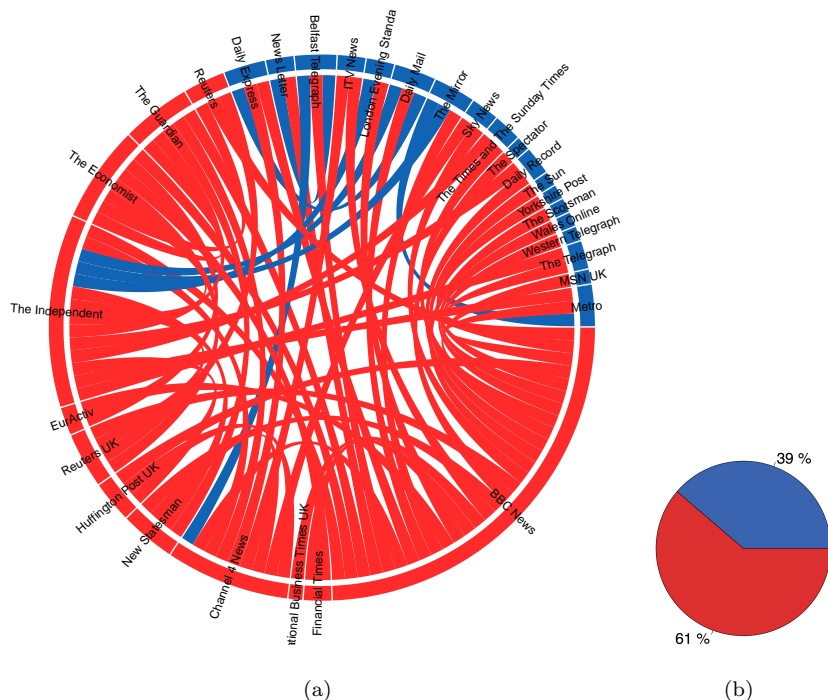


Figure 1: Backbone structure for the bipartite projection of the pages-users graph G_p (a) and percentage of pages in the different communities (b). Colors indicate the membership of users in the different communities (*blue* for C1, *red* for C2) detected by the FG algorithm, while for the extraction of the backbone we considered the level of significance $\alpha = 0.03$.

Taking into account the positive meaning of the like as a feedback to a post, we characterize how contents from the two communities detected in G_p are consumed by Facebook users. We define the users polarization by likes (reps., comments) as

$$\varrho(u) = (y - x)/(y + x),$$

where y is the number of likes (resp., comments) that user u left on posts of C1 and x the number of likes (resp., comments) left on posts of C2. Thus, a user u is said to be polarized towards C1 (resp., C2) if $\varrho(u) = 1$ (resp., -1). In Fig. 2 we report the Probability Density Function (PDF) of users polarization by likes (left panel) and comments (right panel). We find that $\varrho(u)$ is sharply bimodal in both cases, denoting that the majority of users may be divided into two main




Figure 2: Probability density function (PDF) for the users polarization $\varrho(u)$ by likes (*left*) and by comments (*right*). $\varrho(u) = 1$ (resp. $\varrho(u) = -1$) indicates that users u is polarized towards C2 (resp., C1).

groups referring to the two communities of Fig. 1(a). Tab. 2 shows the number of polarized users towards both communities by likes and comments.

	C1	C2
<i>Likes</i>	1, 037, 969	255, 930
<i>Comments</i>	168, 680	75, 851

Table 2: Number of polarized users towards both communities by likes and comments.

160 Thus we have shown that users form two well separated communities. Indeed the Brexit debate shows features very similar to those already observed for other phenomena, such as political bloggers [29] or misinformation. In the latter case, researchers confirmed the existence of echo chambers both on Facebook [5, 6] and YouTube [30] and showed that users interact primarily with either conspiracy-
165 like or scientific pages. However, such a symmetry takes on a particular relevance in this work, because the emergence of the two echo chambers is completely spontaneous and no a priori categorization has been performed.

We now want to compare users' activities on posts from Brexit pages. In Fig. 3(a) we report the Complementary Cumulative Distribution Function (CCDF) of likes (left) and comments (right) made by users polarized towards both communities, while in Fig. 3(b) we report the CCDF of the lifetime of polarized users. The lifetime is defined as the temporal distance, in terms of days, between the first and last comment made by any given user. We fitted the distributions in Fig. 3(a) with different models (the exponential, the power law, and the log-normal) by means of NLS estimation, goodness of fit tests are based on the maximization of the log-likelihood. Then, we pairwise compare the distributions of the number of likes and comments by users polarized towards either community by means of the *Kolmogorov-Smirnov* (KS) test. Results for the best fit and KS test are reported in Tab. 3. We may notice that the distributions are all best fitted by the exponential model, with the exception of that of the number of likes by users from C2, that is best fitted by the power-law. Also, we fail to reject the null hypothesis of equivalence of the two distributions in the case of the number of comments by users from either community.

Distribution	Best Fit	Estimated Parameters
# of likes (C1)	exponential	$\hat{a} = 1.49, \hat{b} = 0.00003$
# of likes (C2)	power-law	$\hat{a} = 1.18, \hat{b} = -0.03$
# of comments (C1)	exponential	$\hat{a} = 1.55, \hat{b} = 0.00005$
# of comments (C2)	exponential	$\hat{a} = 1.58, \hat{b} = 0.0001$

KS test		
Compared Distributions	D (C)	p-value
# of likes (C1/C2)	0.075 (0.004)	2×10^{-16}
# of comments (C1/C2)	0.004 (0.007)	0.349

Table 3: **Fit of distributions** from Fig. 3 and results of Kolmogorov-Smirnov tests.

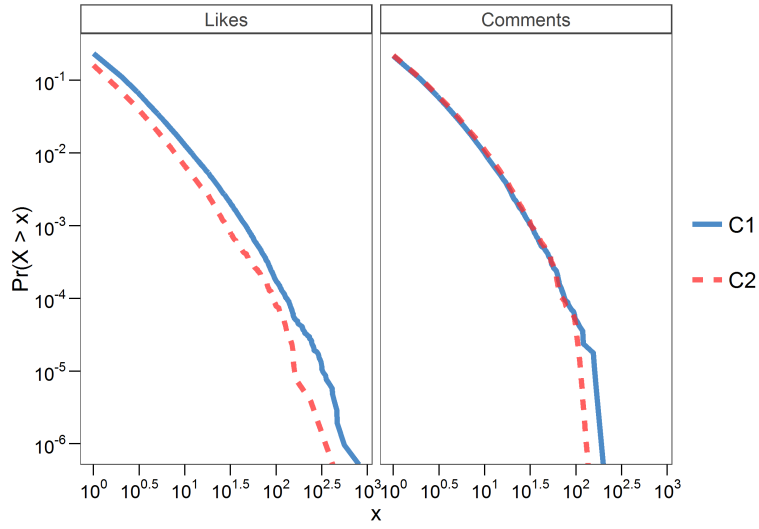
Interestingly, although users tend to focus on contents coming just from one of the two communities, the distributions of their attention patterns are very

similar, and even equal in the case of comments.


2.2. Emotional Dynamics Inside and Between Communities

Our analysis provides evidence of the existence of two well segregated echo chambers: users tend to focus on one narrative and to ignore the other. Such a pattern might be driven by the way in which contents are debated on pages, i.e., is such a way that matches their own users' preferences. To shade light on this aspect, we want to measure the distance among the sentiment of the users w.r.t. the same topic. Thus, we analyze how the subject of a post is presented to the users. To perform the analysis we make use of IBM Watson™ Alchemy-Language service API [31], that allows us to extract semantic meta-data from posts content. Such a procedure applies machine learning and natural language processing techniques aimed to analyze text by automatically extracting relevant entities, their semantic relationship as well as the emotional sentiment they express [32]. In particular, we extract the sentiment and the main concepts presented by each post of the dataset, whether it has a textual description or a link to an external document. The AlchemyAPI tools make use of the language patterns surrounding the input text looking for signals that denote the sentiment and exploring information based on the concepts behind such an input. Thus, a concept is a high-level conceptual association identified in the content provided as input to the service. Input content is auto-tagged against a concept graph, which formally represents the relationships between the concepts contained in the data on which it is based.

Fig. 4 shows the sentiment distribution of posts on both communities. The sentiment score is defined in the range $[-1, 1]$, where -1 is negative, 0 is neutral, and 1 is positive. We may observe a negative overall pattern for both categories, although clearly more pronounced for posts of C1. Notice that we consider how subjects are presented in a post; here we do not take into account the sentiment that the post may elicit in the reader, or the sentiment of users involved in the discussion.



(a)



(b)

Figure 3: (a) CCDF of likes (*left*) and comments (*right*) made by users polarized towards C1 (*solid blue*) and towards C2 (*dashed red*). (b) CCDF of the lifetime of users polarized towards either C1 (*solid blue*) or 2 (*dashed red*). The lifetime is computed as the temporal distance, in terms of days, between the first and last comment made by any given user.




Figure 4: Probability Density Function (PDF) of posts sentiment score on C1 (*solid blue*) and C2 (*dashed red*). The sentiment score is defined in the range $[-1, 1]$, where -1 is negative, 0 is neutral, and 1 is positive

215 *Controversial Concepts: Emotional Distance and Users' Response*


We now want to understand how users of the two echo chambers perceive the issues debated on their pages. Thus, we focus on the concepts of each echo chamber: 102 such concepts are shared by both communities, for a total of 1,520 posts ($1,258 \in C1$, $262 \in C2$) and 115,958 comments ($95,357 \in C1$,
 220 $20,601 \in C2$). For each concept we compute its average sentiment – i.e., the mean of the sentiment of all the posts where it appears. The emotional distance between two concepts is defined as the difference between the average sentiment of the concept in C2 and that in C1. Since we are interested in identifying the most controversial concepts, we consider only those concepts for which the
 225 emotional distance (in absolute value) between the two communities is greater than 0.2. Fig. 5 shows, for each concept, the emotional distance between the two echo chambers. More specifically, the top panel (a) of Fig. 5 includes the 52 concepts that are presented in a more negative way in community C1 w.r.t. C2, while the bottom panel (b) includes the 48 concepts that are presented in
 230 a more negative way in community C2 w.r.t. C1. In both panels concepts are

shown in descending order by the largest to the smallest emotional distance. Thus, concepts on the left are those discussed with the greatest difference in sentiment, while those on the right are discussed in a much more similar way by both echo chambers.


235 So far we have analyzed how subjects are debated in the posts of Brexit pages. What about the emotional response of users to such posts? To this aim we take all the comments (115,958) of posts including one of the concepts and compute their sentiment score through AlchemyAPI. Thus, to each comment is associated a sentiment score in $[-1,1]$, where -1 is negative, 0 is neutral,
240 and 1 is positive. For each post (resp., user) we compute the average sentiment of its (resp., her) comments – i.e., the mean of the sentiment of all comments on the post (resp., made by the user). Then, for each concept, we consider the emotional distance between the average sentiment of the post and that of its users. Fig. 6 shows the emotional response of users to posts of C1 (a)
245 and C2 (b) debating one of the listed controversial topics. Only concepts for which the emotional distance (in absolute value) between the two communities is greater than 0.2 have been taken into account. In both panels a vertical dashed line denotes a change in users’ response: concepts on the left are those for which users’ response is more negative than the sentiment expressed in the
250 post, and vice versa for those on the right. We may notice that users tend to react negatively to the content of the posts, independently of their reference community.

Conclusions

We address the online discussion around Brexit on Facebook by means of
255 a quantitative analysis on a sample of 5K posts from 38 pages linked to official UK news sources. We observe the spontaneous emergence of two separate communities, where the connections among pages are the direct result of users’ activity and no reference to the shared contents is implied. We further explore the dynamics of the discussion by looking at the polarization of users from the

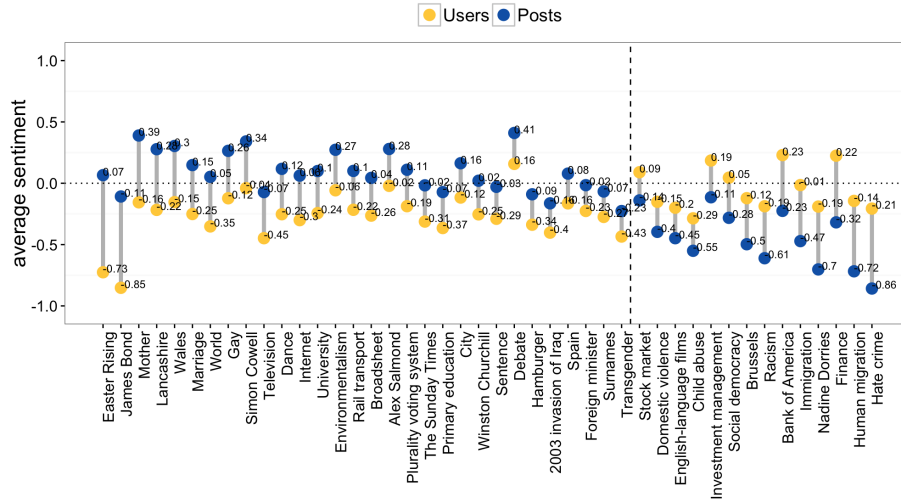


(a)

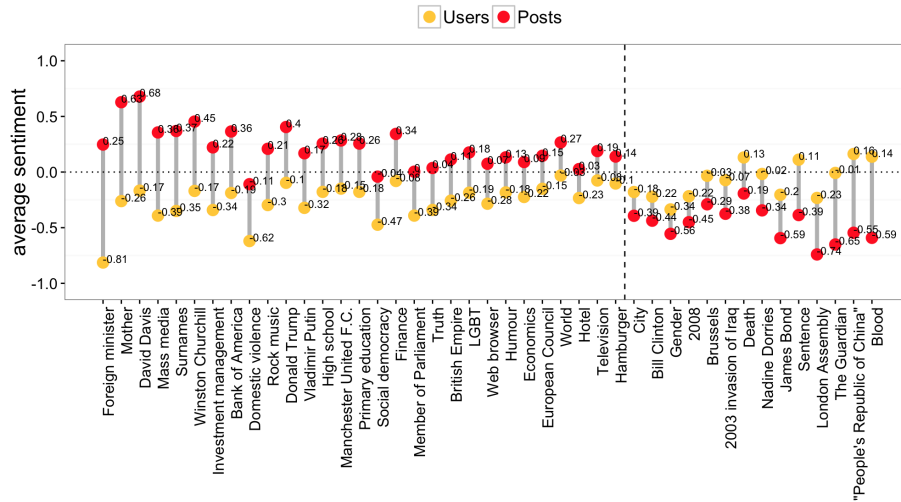


(b)

Figure 5: **Emotional Distance Between Communities.** Emotional distance – i.e., the distance between the average sentiment of a concept on both communities – for each concept debated in both communities. Panel (a) includes the 52 concepts that are presented in a more negative way in community C1 (blue dots) w.r.t. C2 (red dots), while panel (b) includes the 48 concepts that are presented in a more negative way in community C2 (red dots) w.r.t. C1 (blue dots). Concepts are shown in a descending order by the largest to the smallest emotional distance.



(a)



(b)

Figure 6: **Emotional Response to Controversial Concepts.** Panel (a) shows the emotional response of users (yellow dots) to posts of C1 (blue dots) debating one of the listed controversial concepts, while panel (b) shows the emotional response of users (yellow dots) to posts of C2 (red dots). Only concepts for which the emotional distance between the two communities is greater than 0.2 are reported. The vertical dashed lines denote a change in users' response.

260 two communities and their attention patterns. We find a sharply bimodal dis-
tribution for the polarization of users. Users segregation might be driven by
the match between their personal preferences and the way in which contents
are presented. We identify how concepts get received and shape the narrative
inside the echo chamber by measuring both the distance between the sentiment
265 of users w.r.t. the same topic and that of users w.r.t. the “presentation” of the
topic. Firstly, we characterize the structural properties of the discussion by ob-
serving the spontaneous emergence of two well-separated communities; indeed,
connections among pages are the direct result of users’ activity, and we do not
perform any categorization of contents a priori. Then, we explore the dynamics
270 behind discussion: looking at users polarization towards the two communities
and at their attention patterns, we find a sharply bimodal distribution, showing
that users are divided into two main distinct groups and confine their atten-
tion on specific pages. Finally, to better characterize inner group dynamics,
we introduce a new technique which combines automatic topic extraction and
275 sentiment analysis. We compare how the same topics are presented on posts
and the related comments, finding significant differences in both echo chambers
and that polarization reflects on the perception of topics. We first measure the
distance between how a certain concept is presented on the posts by both echo
chambers and then we measure the emotional response of users to such contro-
280 versial topics. Our new measures could be of great interest to identify the most
crucial topics in online debates. Indeed, it is highly likely that the greater the
emotional distance between the same concept in two echo chambers, the greater
users’ polarization. Our results provide important insights for identifying the
determinants of polarization and evolution of the core narratives behind online
285 debating.

References

- [1] N. Newman, D. A. Levy, R. K. Nielsen, Reuters institute digital news report 2015, Available at SSRN 2619576.

- 290 [2] M. Del Vicario, A. Bessi, F. Zollo, F. Petroni, A. Scala, G. Caldarelli,
H. E. Stanley, W. Quattrociocchi, The spreading of misinformation online,
Proceedings of the National Academy of Sciences 113 (3) (2016) 554–559.
- [3] W. Quattrociocchi, A. Scala, C. R. Sunstein, Echo chambers on facebook,
Available at SSRN.
- 295 [4] A. Bessi, F. Petroni, M. Del Vicario, F. Zollo, A. Anagnostopoulos,
A. Scala, G. Caldarelli, W. Quattrociocchi, Viral misinformation: The role
of homophily and polarization, in: Proceedings of the 24th International
Conference on World Wide Web Companion, International World Wide
Web Conferences Steering Committee, 2015, pp. 355–356.
- 300 [5] F. Zollo, A. Bessi, M. Del Vicario, A. Scala, G. Caldarelli, L. Shekht-
man, S. Havlin, W. Quattrociocchi, Debunking in a world of tribes, arXiv
preprint arXiv:1510.04267.
- [6] A. Bessi, M. Coletto, G. A. Davidescu, A. Scala, G. Caldarelli, W. Quat-
trociocchi, Science vs conspiracy: collective narratives in the age of misin-
formation, PloS one 10 (2) (2015) 02.
- 305 [7] F. Zollo, P. K. Novak, M. Del Vicario, A. Bessi, I. Mozetič, A. Scala,
G. Caldarelli, W. Quattrociocchi, Emotional dynamics in the age of misin-
formation, PloS one 10 (9) (2015) e0138740.
- [8] C. R. Sunstein, The law of group polarization, Journal of political philos-
ophy 10 (2) (2002) 175–195.
- 310 [9] D. Centola, The spread of behavior in an online social network experiment,
Science 329 (5996) (2010) 1194–1197.
- [10] J. H. Fowler, N. A. Christakis, Cooperative behavior cascades in human
social networks, Proceedings of the National Academy of Sciences 107 (12)
(2010) 5334–5338.

- 315 [11] W. Quattrociocchi, G. Caldarelli, A. Scala, Opinion dynamics on inter-
acting networks: media competition and social influence, *Scientific reports*
4.
- [12] M. J. Salganik, P. S. Dodds, D. J. Watts, Experimental study of inequality
and unpredictability in an artificial cultural market, *Science* 311 (5762)
320 (2006) 854–856.
- [13] A. D. I. Kramer, J. E. Guillory, J. T. Hancock, Experimental evidence of
massive-scale emotional contagion through social networks, *Proceedings of*
the National Academy of Sciences (2014) 8788–8790.
- [14] L. Backstrom, J. Kleinberg, L. Lee, C. Danescu-Niculescu-Mizil, Character-
izing and curating conversation threads: expansion, focus, volume, re-entry,
325 in: *Proceedings of the sixth ACM international conference on Web search*
and data mining, ACM, 2013, pp. 13–22.
- [15] Q. Diao, J. Jiang, F. Zhu, E.-P. Lim, Finding bursty topics from microblogs,
in: *Proceedings of the 50th Annual Meeting of the Association for Compu-*
tational Linguistics: Long Papers-Volume 1, Association for Computational
330 Linguistics, 2012, pp. 536–544.
- [16] D. M. Romero, B. Meeder, J. Kleinberg, Differences in the mechanics of
information diffusion across topics: idioms, political hashtags, and complex
contagion on twitter, in: *Proceedings of the 20th international conference*
on World wide web, ACM, 2011, pp. 695–704.
335
- [17] B. Stolz, H. Harrington, M. A. Porter, The topological 'shape' of brexit,
arXiv preprint arXiv:1610.00752.
- [18] P. N. Howard, B. Kollanyi, Bots, # strongerin, and # brexit: Computational
propaganda during the uk-eu referendum, Available at SSRN 2798311.
- 340 [19] M. Bianchetti, D. Galli, C. Ricci, A. Salvatori, M. Scaringi, Brexit or bre-
main? evidence from bubble analysis, *Evidence from Bubble Analysis* (June
20, 2016).

- [20] R. Steinberger, B. Pouliquen, E. Van Der Goot, An introduction to the europe media monitor family of applications, arXiv preprint arXiv:1309.5290.
- 345 [21] Facebook, Using the graph api, Website, last checked: 19.01.2014 (8 2013).
URL <https://developers.facebook.com/docs/graph-api/using-graph-api/>
- [22] A. Clauset, M. E. Newman, C. Moore, Finding community structure in very large networks, *Physical review E* 70 (6) (2004) 066111.
- 350 [23] P. Pons, M. Latapy, Computing communities in large networks using random walks., *J. Graph Algorithms Appl.* 10 (2) (2006) 191–218.
- [24] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, E. Lefebvre, Fast unfolding of communities in large networks, *Journal of statistical mechanics: theory and experiment* 2008 (10) (2008) P10008.
- 355 [25] U. N. Raghavan, R. Albert, S. Kumara, Near linear time algorithm to detect community structures in large-scale networks, *Physical review E* 76 (3) (2007) 036106.
- [26] W. M. Rand, Objective criteria for the evaluation of clustering methods, *Journal of the American Statistical association* 66 (336) (1971) 846–850.
- 360 [27] L. Hubert, P. Arabie, Comparing partitions, *Journal of classification* 2 (1) (1985) 193–218.
- [28] M. A. Serrano, M. Boguna, A. Vespignani, Extracting the multiscale backbone of complex weighted networks, *Proceedings of the National Academy of Sciences* 106(16) (2009) 6483–6488.
365 URL <http://dx.doi.org/10.1073/pnas.0808904106>
- [29] L. A. Adamic, N. Glance, The political blogosphere and the 2004 us election: divided they blog, in: *Proceedings of the 3rd international workshop on Link discovery*, ACM, 2005, pp. 36–43.

- [30] A. Bessi, F. Zollo, M. Del Vicario, M. Puliga, A. Scala, G. Caldarelli,
 370 B. Uzzi, W. Quattrociocchi, Users polarization on facebook and youtube,
 PloS one 11 (8) (2016) e0159641.
- [31] Alchemy API — Powering the New AI Economy, An IBM Company (July
 2016).
 URL <http://www.alchemyapi.com/>
- 375 [32] A. Gangemi, A comparison of knowledge extraction tools for the semantic
 web, in: Extended Semantic Web Conference, Springer, 2013, pp. 351–366.

Appendix A

We analyze the community structure of the Brexit pages graph by applying
 four different community detection algorithms: Fast Greedy (FG) [22], Walktrap
 380 (WT) [23], Multilevel (ML) [24], and Label Propagation (LP) [25]. The Brexit
 pages graph is the bipartite projection of the pages-users graph G_p where nodes
 are Brexit pages and two pages are connected if at least one user liked a post
 from each of them. The weight of a link is determined by the number of users
 in common between the two pages. To compare the various community struc-
 385 tures we use the Rand method that computes the similarity between different
 clustering methods by considering how nodes are assigned by each community
 detection algorithm [26, 27]. Results for the comparison are reported in Tab. 4.

	FG	WT	ML	LP
FG	1	0.69	0.90	0.50
WT	–	1	0.63	0.32
ML	–	–	1	0.52
LP	–	–	–	1

Table 4: The Rand similarity index among the partitions obtained by different community detection algorithms.

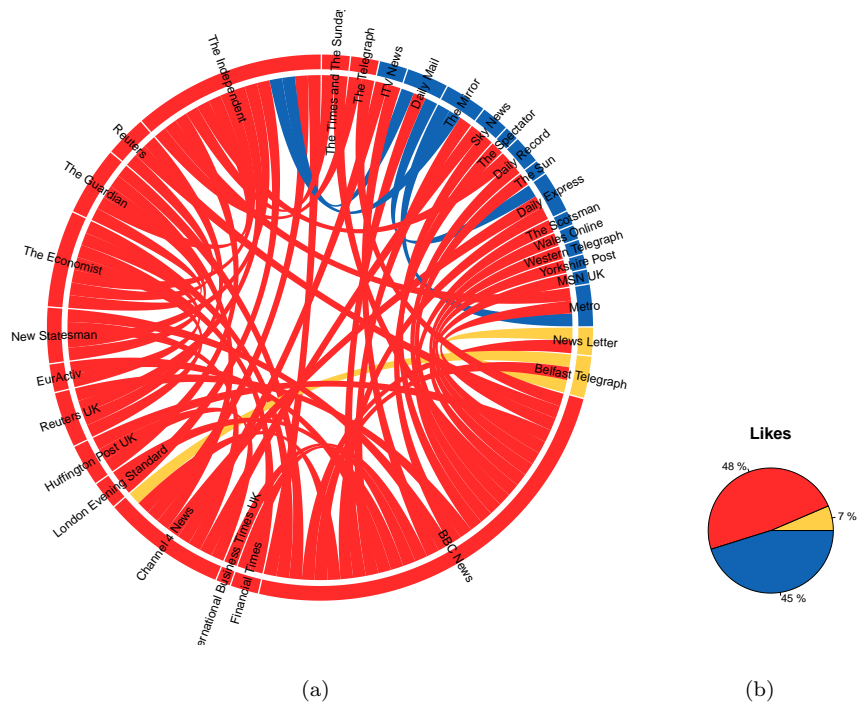


Figure 7: Backbone structure for the bipartite projection of the pages-users graph G_p (a) and percentage of pages in the different communities (b). Colors indicate the membership of users in the different communities (*blue* for C1, *red* for C2, and *yellow* for C3) detected by the WT algorithm, while for the extraction of the backbone we considered the level of significance $\alpha = 0.03$.

In Fig. 7(a) we show the backbone structure of G_p . Colors (resp., blue, red, and yellow) represent the membership to one of the three communities (resp., C1, C2, and C3) detected by the Walktrap (WT) algorithm (see *Methods* section for further details). Fig. 7(b) reports the percentage of pages in each community. Differently from the FG algorithm, WT detects three main communities and five isolated nodes. However, the 27% of nodes belongs to the same community in both cases.

Our analysis underlines the spontaneous emergence of separate communities active on Brexit pages, where the connections among pages are a simple result of the interaction of users on them.

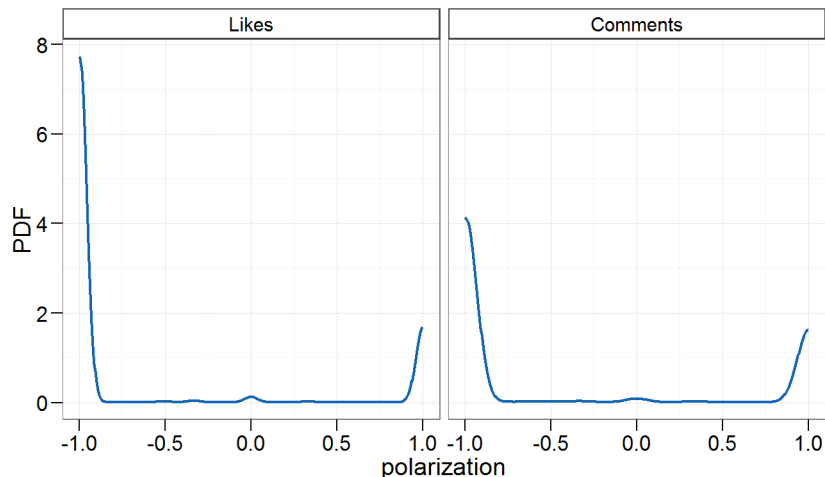


Figure 8: Probability density function (PDF) for the users polarization $\varrho(u)$ by likes (*left*) and by comments (*right*). $\varrho(u) = 1$ (resp. $\varrho(u) = -1$) indicates that users u is polarized towards C2 (resp., C1).

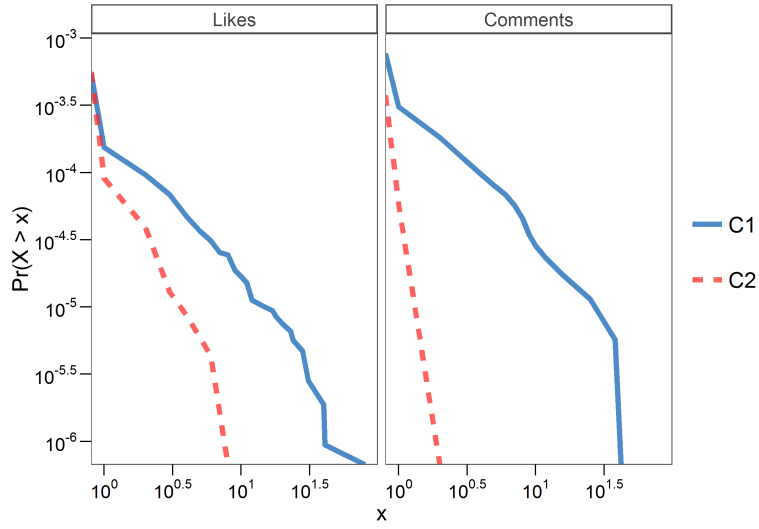
Taking into account the positive meaning of the like as a feedback to a post, we characterize how contents from the two biggest communities detected in G_p are consumed by Facebook users. We define the users polarization by likes (reps., comments) as

$$\varrho(u) = (y - x)/(y + x),$$


where y is the number of likes (resp., comments) that user u left on posts of C2 and x the number of likes (resp., comments) left on posts of C1. Thus, a user u is said to be polarized towards C2 (resp., C1) if $\varrho(u) = 1$ (resp., -1). In Fig. 8 we report the Probability Density Function (PDF) of users polarization by likes (left panel) and comments (right panel). We find that $\varrho(u)$ is sharply bimodal in both cases, denoting that the majority of users may be divided into two main groups referring to the two communities of Fig. 7(a).

Thus, we have shown that users form well segregated communities. We now want to compare their activities on posts from Brexit pages. In Fig. 9(a) we report the Complementary Cumulative Distribution Function (CCDF) of likes

(left) and comments (right) made by users polarized towards both communities, while in Fig. 9(b) we report the CCDF of the lifetime of polarized users. The
410 lifetime is defined as the temporal distance, in terms of days, between the first and last comment made by any given user.



(a)



(b)

Figure 9: (a) CCDF of likes (*left*) and comments (*right*) made by users polarized towards C1 (*solid blue*) and towards C2 (*dashed red*). (b) CCDF of the lifetime of users polarized towards either C1 (*solid blue*) or 2 (*dashed red*). The lifetime is computed as the temporal distance, in terms of days, between the first and last comment made by any given user.

Appendix B

In this section we provide the list of all Facebook pages of news sources whose legal head office (at least one of them) is located in the United Kingdom.

415 Pages engaged in the debate around Brexit are denoted by Y (N, otherwise), followed by the community to which they belong (C1 or C2).

ID	Page Name	Facebook Code	Brexit
1	BBC News	228735667216	Y(C1)
2	Channel 4 News	6622931938	Y(C1)
3	Euractiv (English)	15299247059	Y(C1)
4	Financial Times	8860325749	Y(C1)
5	Huffingtonpost UK	143753582359049	Y(C1)
6	International Business Times UK	224377357631653	Y(C1)
7	New Economics Foundation	110275553302	Y(C1)
8	New Statesman	100959719644	Y(C1)
9	Open Europe Today	321253057971308	Y(C1)
10	Reuters	114050161948682	Y(C1)
11	Reuters UK	208314602512037	Y(C1)
12	The Economist	6013004059	Y(C1)
13	The Guardian UK	10513336322	Y(C1)
14	The Independent	13312631635	Y(C1)
15	The Register	206419956048907	Y(C1)
16	WN.com	229101503845879	Y(C1)
17	Belfast Telegraph	237692023818	Y(C2)
18	Daily Express	129617873765147	Y(C2)
19	Daily Mail	164305410295882	Y(C2)
20	Daily Record	187523381277554	Y(C2)
21	East Anglian Daily Times	6478299951	Y(C2)
22	ITV News	148007467671	Y(C2)
23	London Evening Standard	165348596842143	Y(C2)
24	Metro (UK)	117118184990145	Y(C2)
25	MSN UK	358837740527	Y(C2)
26	News Letter	117370764948881	Y(C2)

27	Nottingham Post	309833935716287	Y(C2)
28	Sky News	164665060214766	Y(C2)
29	The Mirror	6149699161	Y(C2)
30	The Scotsman	293226174987	Y(C2)
31	The Spectator	111263798903232	Y(C2)
32	The Sun	161385360554578	Y(C2)
33	The Times and Sunday Times	147384458624178	Y(C2)
34	Wales Online	21226447182	Y(C2)
35	Wandsworth Guardian	113349742029506	Y(C2)
36	Western Telegraph	180521675319022	Y(C2)
37	Yorkshire Post	316795048375439	Y(C2)
38	Airforce Technology Website	376588539031515	N
39	Azo Mining	195005930530874	N
40	BBC Radio	1470145583204820	N
41	Cafebabel (English)	357343795001	N
42	City A.M.	213682385348579	N
43	Dunmow Broadcast	181182540669	N
44	EU business	215108901846669	N
45	Euromoney	192279900885723	N
46	European Railway Review	404359882930504	N
47	Expatica	206982432584	N
48	Farming Life	243070359106664	N
49	FCO - Foreign and Commonwealth Office	408582579294175	N
50	Harborough Mail	219817851378553	N
51	Herald Scotland	271154343382	N
52	Inmarsat	317156988374684	N
53	Lydian International	186900121339682	N
54	Mining Technology	326019370778750	N
55	Mondo Visione	169767016460715	N
56	MoneyWeek	110326662354766	N
57	Monsters and Critics	193326863118	N
58	New Civil Engineer	166793706822441	N
59	OneWorld.net: Palestine	106968052697581	N

60	Oxford Analytica	160525917321265	N
61	Pan European Networks	230201663697109	N
62	Publish What You Pay	176624229034172	N
63	Railway Magazine	135345903226042	N
64	Routes News	126251777434574	N
65	Seatrade Global	470795739645931	N
66	Survival International	19668531552	N
67	Tax-News	375456009146619	N
68	The Argus	57197526698	N
69	The Courier	325681791214	N
70	The International Institute for Strategic Studies	29840385993	N
71	The Scottish Government	200786289976224	N
72	The Telegraph	143666524748	N
73	The Visitor	68554461041	N
74	This is Africa	779213412106756	N
75	This is Derbyshire	142370589115824	N
76	This is Staffordshire	11878899813	N
77	Thomson Reuters Foundation	31301735406	N
78	World Fishing - The Magazine for Fishing	552321618120006	N
79	Cyprus Expat News	357342727764507	N
80	African Business Magazine	114117578656259	N
81	African Review	507239115959583	N

Table 5: **UK Facebook News Sources and Brexit Community Membership:** List of all Facebook pages of news sources whose legal head office (at least one of them) is located in the United Kingdom. Pages engaged in the debate around Brexit are denoted by *Y*, followed by the community to which they belong.