

APPROVED: 20 April 2019

doi: 10.2903/j.efsa.2019.e170720

Dealing with digital misinformation: a polarised context of narratives and tribes

Fabiana Zollo

Ca' Foscari University of Venice, Italy and Centre for the Humanities and Social Change, Venice, Italy

Abstract

The advent of the internet and social networks has revolutionised the information space and changed the way in which we communicate and get informed. On the internet, a huge amount of information competes for our (limited) attention. Moreover, despite the increasing quantity of contents, quality may be poor, making the environment particularly florid for misinformation spreading. In such a context, our cognitive biases emerge, first and foremost, confirmation bias, i.e. the human tendency to look for information that is already in agreement with one's system of beliefs. To shed light on the phenomenon, we present a collection of works investigating how information gets consumed and shapes communities on Facebook. We find that confirmation bias plays a crucial role in content selection and diffusion, and we provide empirical evidence of the existence of echo chambers, i.e. well separated and polarised groups of like-minded users sharing a same narrative. Immersed in these bubbles, users keep framing and reinforcing their world view, ignoring information dissenting from their preferred narrative. In this scenario, corrections in the form of fact-checking or debunking attempts seem to fail and have instead a backfire effect. To contrast misinformation, smoothing polarisation is so essential, and may require the design of tailored counter-narratives and appropriate communication strategies, particularly for sensitive topics.

© 2019 European Food Safety Authority. *EFSA Journal* published by John Wiley and Sons Ltd on behalf of European Food Safety Authority.

Keywords: misinformation, echo chambers, polarisation, narratives, fake news, confirmation bias

Correspondence: fabiana.zollo@unive.it

Acknowledgements: This article is based on coauthored material. The author wishes to thank Walter Quattrociochi, Antonio Scala, Ana Lucia Schmidt, and all other coauthors.

Suggested citation: Zollo F, 2019. Dealing with digital misinformation: a polarised context of narratives and tribes. *EFSA Journal* 2019;17(S1):e170720, 15 pp. <https://doi.org/10.2903/j.efsa.2019.e170720>

ISSN: 1831-4732

© 2019 European Food Safety Authority. *EFSA Journal* published by John Wiley and Sons Ltd on behalf of European Food Safety Authority.

This is an open access article under the terms of the [Creative Commons Attribution-NoDerivs License](#), which permits use and distribution in any medium, provided the original work is properly cited and no modifications or adaptations are made.

Reproduction of the images listed below is prohibited and permission must be sought directly from the copyright holder:

Figure 2: © PLoS ONE; Figure 3: © PLoS ONE; Figure 4: © PLoS ONE; Figure 5: © Social Networks, Elsevier; Figure 6: © Social Networks, Elsevier; Figure 7: © Vaccine, Elsevier; Figure 8: © PNAS



The EFSA Journal is a publication of the European Food Safety Authority, an agency of the European Union.



Table of contents

Abstract.....	1
1. Introduction.....	4
2. Data and methodologies.....	5
2.1. Ethics statement.....	5
2.2. Misinformation.....	5
2.3. Narratives and news consumption.....	5
2.3.1. Brexit.....	5
2.3.2. Vaccines.....	5
2.3.3. Anglophone news on a global scale.....	6
3. Results and discussion.....	6
3.1. Echo chambers.....	7
3.2. The pseudoscientific narrative.....	7
3.3. Emotional dynamics between tribes.....	8
3.4. Response to dissenting information.....	9
3.5. Narratives.....	10
3.5.1. Brexit.....	10
3.5.2. Vaccines.....	12
3.6. News consumption and selective exposure.....	12
References.....	14
Abbreviations.....	15

1. Introduction

In her speech at 2016 Yale College Class Day, the US Ambassador Samantha Power encouraged the audience to break out of their own echo chambers and strive to engage with people they disagreed with, this being the only way to not only challenge their view, but also to acknowledge that they might be right.¹ The echo chambers thesis is nothing new: there is plenty of evidence of their presence on the internet (O'Hara and Stevens, 2015), and their effect on opinion formation (Sunstein, 2012).

Over the years, the internet has been celebrated for sharing information, disseminating knowledge, promoting freedom, debate and – not least – the emergence of the so-called collective intelligence (Lévy and Bononno, 1997). Indeed, the advent of the internet – and social networks – has revolutionised the information space by allowing anyone to post opinions and produce online content. Traditional media such as print, radio and television have been coupled with a heterogeneous mass of alternative news sources, in which information is no longer mediated as it used to be. Disintermediation is not the only change, it comes with a hyperconnected world where communicating with other people is extremely easy and rapid, without any temporal or spatial barrier. The small-world theory tested by Milgram in the 1960s (Travers and Milgram, 1977) resulted in the so-called idea of six degrees of separation, being six the number of intermediaries necessary to connect two individuals. In 2016, researchers at Facebook found out the same number to be 3.57, meaning that today each person is connected to every other person by an average of three and a half other people (Bhagat et al., 2016).

That is a simple example of how the advent of the new technologies changed the way in which we communicate and get informed. We are increasingly relying on the internet – and social media – in both our civic and daily life. As of the fourth quarter of 2018, Facebook had 2.32 billion monthly active users,² while WhatsApp use for news has almost tripled since 2014 (Nic et al., 2018). Indeed, the majority of these users prefers to access news through social media, search engines or news aggregators, rather than going directly to a news website. At the same time, smartphone reach for news is significant, and might affect the way we consume information (and the time that we devote to processing it). On the internet, a huge amount of information competes for our attention, which is limited. Moreover, despite the increasing quantity of contents, quality may be poor, due to issues of content monetisation and the persisting reduction of investments in the news production and distribution (AGCOM, 2018).

In such a context, it is difficult to implement our abilities to analyse, reflect and draw conclusions. Instead, our cognitive biases emerge. As human beings, we need such biases to interpret the reality. They are shortcuts, heuristics that we use to simplify the reality and (re)act rapidly. Unfortunately, such a simplification may come with undesired effects. While these cognitive mechanisms are often fundamental to our survival, they might also act as mental traps and mislead us. Among these, confirmation bias is the human tendency to look for information that is already in agreement with one's system of beliefs. We will see later that, despite the availability of a huge (if not infinite) variety of information, online users tend to fragment into bubbles, each one with its own narrative and world view, the previously mentioned echo chambers (Zollo and Quattrocchi, 2018a,b). The combination of confirmation bias with the hyperconnected information space makes the environment florid for misinformation. Since 2013, the World Economic Forum (WEF) has been placing the global risk of massive digital misinformation at the core of technological and geopolitical risks such as rising religious fanaticism, cyber-attacks and terrorism (Howell, 2013). Misinformation, disinformation and fake news are dangerous because they might influence public opinion and have an impact on the real world. Examples in this regard are numerous, from the death of two men over excessive salt consumption to prevent the Ebola virus according to a viral message online,³ to Texas Governor Abbott's decision to alert the State Guard to monitor the Jade Helm 15 operation, to the increasing spread of HIV/AIDS due to denialism beliefs.

To shed light on the phenomenon, in the following sections we will discuss the role of confirmation bias in (mis)information consumption in the digital era. To this aim, we will adopt a cross-methodological approach to formulate data-driven models and quantitatively analyse social dynamics. In this direction, we can benefit from the large availability of data from online social networks and provide empirical evidence of the existence of echo chambers that has been so far dependent on surveys (and so self-reporting) or based on small/inadequate samples (O'Hara and Stevens, 2015).

¹ Source: <https://news.yale.edu/2016/05/22/ambassador-samantha-powers-2016-class-day-address>

² Source: Statista 2019.

³ Source: <https://www.vanguardngr.com/2014/08/ebola-two-die-drinking-salt-water-jos/>

The manuscript is structured as follows: in Section 2, data and methodologies are introduced; in Section 3, results are presented and discussed.

2. Data and methodologies

As anticipated, our approach takes advantage of both the question-framing capabilities of social sciences and the experimental and quantitative tools of hard sciences. In this section, we present a collection of data sets (Bessi et al., 2015, 2016; Zollo et al., 2015, 2017; Del Vicario et al., 2016, 2017; Schmidt et al., 2017, 2018) that lay the groundwork for our analysis.

2.1. Ethics statement

The entire data collection process was performed exclusively by means of the Facebook Graph API,⁴ which is publicly available and can be used through one's personal Facebook user account. We used only publicly available data (users with privacy restrictions are not included in our data set). Data were downloaded from Facebook pages that are public entities. Users' content contributing to such entities is also public unless users' privacy settings specify otherwise and, in that case, it is not available to us. When allowed by users' privacy specifications, we accessed public personal information. However, in our study we used fully anonymised and aggregated data. We abided by the terms, conditions and privacy policies of Facebook.

2.2. Misinformation

We define two distinct and opposite narratives, i.e. (a) 'conspiracy'; and (b) 'scientific' information sources (Bessi et al., 2016; Del Vicario et al., 2016). Notice that we did not focus on the quality or veracity of their contents, but rather on their verifiability. Indeed, while producers of scientific information as well as data, methods and outcomes are readily identifiable and available, the origins of conspiracy and pseudoscientific theories are often unknown and their content is strongly disengaged from mainstream society, sharply diverging from recommended practices. In that spirit, we identified two main categories of Facebook pages: the first category (Conspiracy) includes all pages diffusing conspiracy information, i.e. pages that disseminate controversial information, most often lacking supporting evidence and sometimes contradictory of the official news. The second category is that of scientific dissemination and includes institutions, organisations, scientific press, having the main mission to diffuse scientific knowledge. Finally, we identified an additional category of pages, that of 'debunking', i.e. information aiming at correcting false conspiracy theories and untruthful rumours circulating online (Zollo et al., 2017).

To produce our data sets, we categorised conspiracy and scientific news sources available on Facebook with the help of several experts active in debunking fake news and conspiracy theories.⁵ To validate the list, all pages were then manually checked by looking at their self-description and the type of promoted content. The exact breakdowns of the Italian and US Facebook data sets are reported in Tables 1 and 2, respectively.

2.3. Narratives and news consumption

2.3.1. Brexit

Following the Europe Media Monitor (EMM) list (Steinberger et al., 2013), we selected all news outlets (and their related Facebook page) whose legal head office (at least one of them) was located in the United Kingdom (Del Vicario et al., 2017). For each page, we downloaded all the posts from 1 January to 15 July 2016, as well as all the related likes and comments. All UK-based pages were then filtered to include only those pages (38) engaged in the debate around Brexit, i.e. that posted at least one news story about the Brexit. The exact breakdown of data is provided in Table 3.

2.3.2. Vaccines

The data set was generated through requests to Facebook for pages containing the keywords 'vaccine', 'vaccines' or 'vaccination' in their name or description (Schmidt et al., 2018). We then filtered

⁴ Please see: <https://developers.facebook.com/docs/graph-api/>

⁵ We are grateful to *Skepti Forum*, *Skeptical spectacles*, *Butac*, and *Protesi di Complotto*.

the raw Facebook data to include only the ones relevant for the study. Inclusion criteria were language (English), a minimum level of activity (at least 10 posts), date of the posts (between 1 January 2010 to 31 May 2017) and relationship of the page to vaccination. This last step was essential, as having one of the keywords in the description does not necessarily mean the page's topic is about vaccines. Some examples of those false-positive search results are the pages The Vaccines (an UK music band), and Arthur D'Vaccine (a comedian). From the resulting set of Facebook pages, we downloaded all posts as well as all the likes and comments made on those posts. Considering the content of the posts made on the pages, all the Facebook pages were also manually classified by two raters into two groups: 'pro-vaccines' and 'anti-vaccines'. The exact breakdown of data is provided in Table 3.

2.3.3. Anglophone news on a global scale

Starting from the list provided by EMM (Steinberger et al., 2013), we selected all the Anglophone news sources and their related page on Facebook (Schmidt et al., 2017). The downloaded data from each page included all the posts made from 1 January 2010 to 31 December 2015, as well as all the likes and comments on such posts. The EMM list also includes the country and the region of each news source. For an accurate mapping on the globe, we also collected the geographical location (latitude and longitude) of each page. A breakdown of data is reported in Table 3.

Table 1: Breakdown of the Italian Facebook data set

Facebook IT	Science	Conspiracy
Pages	34	39
Posts	62,705	208,591
Likes	2,505,399	6,659,382
Comments	180,918	836,591
Likers	332,357	864,047
Commenters	53,438	226,534

Table 2: Breakdown of the US Facebook data set

Facebook US	Science	Conspiracy	Debunking
Pages	83	330	66
Posts	262,815	369,420	47,780
Likes	453,966,494	145,388,117	3,986,922
Comments	22,093,692	8,304,644	429,204
Likers	39,854,663	19,386,131	702,122
Commenters	7,223,473	3,166,726	118,996

Table 3: Data sets breakdown

Facebook	Brexit (UK)	Vaccines (English)		Anglophone news
		Anti-	Pro-	
Pages	38	98	145	920
Posts	5,039	189,759	108,259	12,825,291
Likes	2,504,956	12,696,440	11,459,295	3,621,383,495
Comments	469,397	1,351,839	749,209	366,406,014
Users	–	1,277,170	1,388,677	376,320,713
Likers	1,365,821	1,145,650	1,325,511	360,303,021
Commenters	259,078	271,598	146,196	60,115,975

3. Results and discussion

Our aim was to analyse social dynamics in a quantitative way using digital traces left by users on online social media. In particular, to study the role of confirmation bias in information spreading, in the following sections we will show the existence of echo chambers on Facebook and (i) analyse users'


behaviour on different and contrasting narratives; (ii) address users' emotional dynamics within and between the communities; and (iii) investigating the response to dissenting information. Furthermore, we will extend our analysis to the cases of Brexit and vaccines, and investigate users' news consumption habits on a global scale.

3.1. Echo chambers

We start our analysis by focusing on how information gets consumed by users on Facebook, considering both the Italian (Bessi et al., 2016; Del Vicario et al., 2016) and the US (Zollo et al., 2017) data sets. At the time of downloading, Facebook's paradigm allowed users to interact with contents by means of likes, comments and shares (reactions were not implemented yet). Each action has a particular meaning (Ellison et al., 2007): While a 'like' represents a positive feedback to the post, a 'share' expresses the desire to increase the visibility of a given information; finally, a 'comment' is the way in which the debate takes form around the topic of the post.

We now want to understand if users' engagement with a specific kind of content can become a good proxy to detect groups of like-minded users, i.e. echo chambers. Assume that a user u has performed n and m likes on scientific and conspiracy posts, respectively, and let $\rho(u) = (m - n) / (m + n)$. So, we say a user u to be polarised towards Science if $(\rho(u) \geq 0.95)$, i.e. she left more than 95% of her likes on scientific posts; vice versa, we say a user u to be polarised towards Conspiracy if $(\rho(u) \leq -0.95)$, i.e. she left more than 95% of her likes on Conspiracy posts. Figure 1 shows the probability density functions (PDFs) of users' polarisation on both the Italian (left) and US (right) Facebook. For both data sets, we may observe a sharply bimodal distribution showing two main peaks by values -1 and 1 . Indeed, the majority of users is polarised either towards Science (~ 1) or Conspiracy (~ -1), eliciting the formation of two well segregated groups of users.

So, our results confirmed the existence of echo chambers. Indeed, distinct narratives aggregate users into separate, polarised communities where they frame and reinforce their world view by acquiring information that confirms their preferred narrative and interact with like-minded people who share the same system of beliefs.



Sources: Bessi et al. (2016); Zollo et al. (2017).

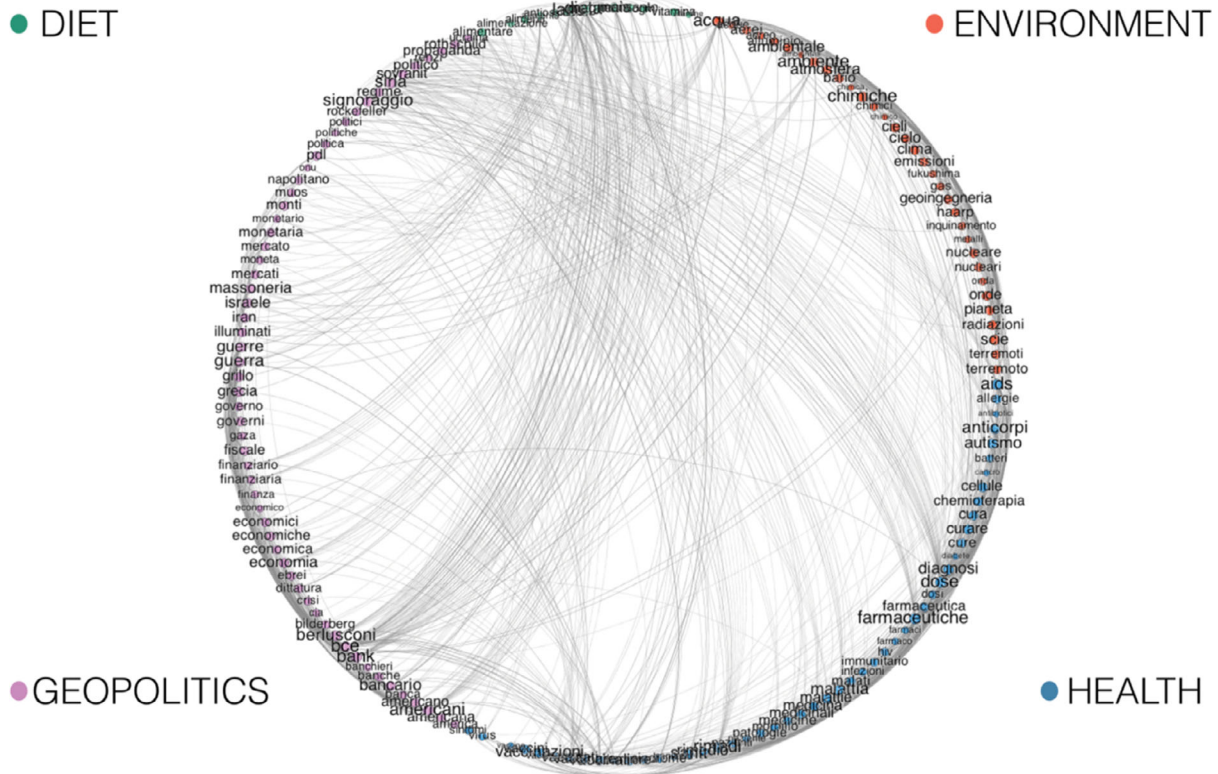
Figure 1: Users' polarisation on Facebook

3.2. The pseudoscientific narrative

Taking a look into the pseudoscientific echo chamber, we may observe that contents mainly refer to four categories: environment, diet, health and geopolitics. Through a semi-automatic topic extraction strategy (Bessi et al., 2015), we were able to identify terms related to the conspiracy storytelling and to derive a co-occurrence network of conspiracy terms, i.e. a graph in which nodes are conspiracy terms, edges bond two nodes if the corresponding terms are found in the same Facebook post, and weights associated to edges indicate the number of times the two terms appear together in the corpus, as shown in Figure 2.

By analysing users' interactions on such posts, we also found that the more a user was active, the more they were likely to span all categories. Within the echo chamber, users jumped from one semantic category to another, and such a probability increased with the user engagement (i.e. number of likes on a single specific category). Each new like on the same category increases, of the 12%, the

probability to pass to a new one. In other words, once involved in a pseudoscientific category, users tend to embrace the overall narrative.




Different colours denote nodes belonging to different semantic categories. In particular, purple nodes belong to geopolitics, red nodes to the environment, blue nodes to health and green nodes to diet. Source: Bessi et al. (2015).

Figure 2: Backbone (Serrano et al., 2009) of conspiracy terms co-occurrence network

3.3. Emotional dynamics between tribes

We now wanted to investigate the emotional dynamics within and across the two echo chambers. In particular, we applied sentiment analysis techniques to the comments of the Italian data set and analysed users' behaviour with respect to scientific and pseudoscientific information (Zollo et al., 2015). The sentiment analysis was based on a supervised machine learning approach in which we first annotated a substantial sample of comments, and then built a classification model, which was applied to associate each comment with one sentiment value: negative (-1), neutral (0), or positive (+1). The sentiment is intended to express the emotional attitude of the user when posting the comment on Facebook.

Looking at how users' sentiment changes with respect to their engagement within the echo chamber, we found that the sentiment of polarised users tend to be more negative than general ones, noticing that the percentage of users whose emotional attitude on Facebook was negative, differed by 11% for Conspiracy, and 8% for Science. We now wondered what happens when such polarised, negative-minded users meet their opponents. To this aim, we selected all the posts commented on at least once by both a user polarised towards Science and a user polarised towards Conspiracy. Slightly more than 25% of the total number of posts met this criterion, so reinforcing the idea that the two communities are strictly separated and do not often interact with one another. Moreover, by analysing how the sentiment changed when the number of comments of the post increased, i.e. when the discussion became longer, we found that the sentiment was increasingly negative, as shown in Figure 3. In other words, the length of the discussion seemed to affect the negativity of the sentiment of the users involved in the debate. In this regard, we may mention that online discussions are often associated with flaming, a hostile and offensive interaction between (or among) users that usually transcends the original topic of discussion (Coe et al., 2014; Zollo and Quattrocchi, 2018b).



Negative (respectively, neutral, positive) sentiment is denoted by red (respectively, yellow, blue) colour. Source: Zollo et al. (2015).

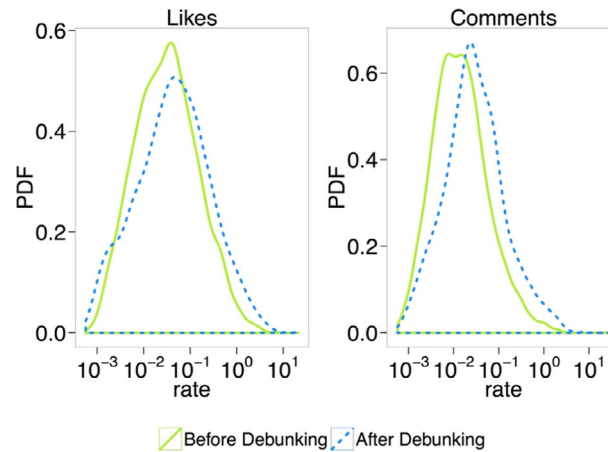
Figure 3: Aggregated sentiment of posts as a function of their number of comments

3.4. Response to dissenting information

Debunkers strive to contrast misinformation spreading by providing fact-checked information to specific topics. As confirmation bias plays a pivotal role in selection criteria, we may assume that debunking activity is likely to sound to Conspiracy users as information dissenting from their preferred narrative. In this section, we investigated the users' response to debunking attempts on US Facebook (Zollo et al., 2017).

Our first question is: Who are the users interacting with debunking content? Looking at their polarisation values, we may notice that the majority of both likes (66.95%) and comments (52.12%) is left by users polarised towards Science, while only a small minority is made by users polarised towards Conspiracy (respectively, 6.54% and 3.88%). This result was interesting, although not surprising: the biggest consumer of debunking information was the scientific echo chamber. Out of 9,790,906 polarised Conspiracy users, 117,736 interacted with debunking posts – i.e. commented a debunking post at least once – just 1.2%. It therefore seems that debunking posts remained confined within the scientific echo chamber, and only a few users usually exposed to unsubstantiated content actively interacted with the correction.

Nevertheless, few users polarised towards Conspiracy did interact with debunking information. We therefore wondered about the impact of such an interaction. Therefore, we compared Conspiracy users' behaviour before and after their first comment to a debunking post. Specifically, Figure 4 shows the liking and commenting rate – i.e. the average number of likes (or comments) to Conspiracy posts per day – before and after the first interaction with debunking. Should the debunking be effective, we would expect that Conspiracy would reduce their engagement with the Conspiracy echo chamber after correction. Instead, what we observed was that their liking and commenting rates on Conspiracy content increased after commenting. This is what is known as a 'backfire effect' (Nyhan and Reifler, 2010), i.e. such users' activity in the Conspiracy echo chamber was reinforced after the interaction, rather than reduced. Paradoxically, conspiracy users not exposed to debunking are 1.76 times more likely to stop interacting with conspiracy contents than those exposed to corrections (Zollo et al., 2017).



Source: Zollo et al. (2017).

Figure 4: Rate, i.e. average number, over time, of likes (left) and comments (right) made to conspiracy posts by users exposed to debunking posts before (solid green line) and after (dashed blue line) the interaction


3.5. Narratives

To better frame the main determinants behind content consumption and the emergence of collective narratives on online social media, in this section we will focus on the discussion around two of the most recent and debated topics: Brexit (Del Vicario et al., 2017) and vaccines (Schmidt et al., 2018).

3.5.1. Brexit

To characterise users' behaviour on Facebook pages engaged in the debate around Brexit, we first analysed their interaction patterns with the pages (Del Vicario et al., 2017). To this aim, we built the page-users graph (Figure 5), in which nodes along the ring are pages, and two pages are connected by an edge if at least one user liked a post from each of them. The weight of an edge was given by the number of users that the two pages had in common. Then, we applied a community detection algorithm⁶ to identify groups of nodes in the network. We were therefore able to detect two main communities, C1 and C2 (respectively, blue and red), observing the spontaneous emergence of two separate communities active around Brexit content, in which connections among pages are a simple result of users' interaction on them. Such a separation is evident in Figure 6, where we report the PDF of users' polarisation taking into account their likes. Indeed, the distribution is sharply bimodal, denoting that the majority of users may be divided into two main groups, i.e. the two communities shown in Figure 5.

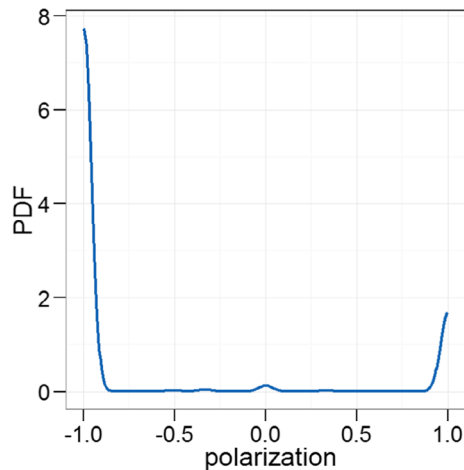
⁶ Communities were detected using the Fast Greedy (FG) algorithm. For more details please refer to (Del Vicario et al., 2017).



Colours denote users' membership (blue for C1, red for C2). Source: Del Vicario et al. (2017).

Figure 5: Community structure around Brexit content

Therefore, the debate around Brexit shows features very similar to those already observed for other political debates (Adamic and Glance, 2005) or misinformation (as shown above). In this specific case, users' polarisation takes on a particular relevance, because the emergence of the two echo chambers was completely spontaneous and no *a priori* categorisation had been performed.




$\rho(u) = 1$ (resp. $\rho(u) = -1$) indicates that users u is polarised towards C2 (resp., C1). Source: Del Vicario et al. (2017).

Figure 6: Brexit. Probability density function (PDF) of users' polarisation $\rho(u)$ considering likes

3.5.2. Vaccines

The internet is pervaded with unsubstantiated claims doubting vaccines safety and efficacy, despite scientific evidence that would prove otherwise. In this section, we describe the evolution of the debate around vaccines on Facebook, taking into account two communities with opposing views, anti-vaccine and pro-vaccine, for a total of 2.6 million users (Schmidt et al., 2018). Even in this case, our findings confirmed the existence of two polarised communities, as shown in Figure 7. Additionally, we found that the two communities displayed different rates of engagement, with the users of the anti-vaccine community being generally more active than those belonging to the pro-vaccine community.

It therefore seems that Facebook allows echo chambers to emerge. In this respect, pro-vaccination campaigns might remain confined to pro-vaccination users, so limiting their efficacy.



$\rho(u) = 1$ (resp. $\rho(u) = -1$) indicates that users u is polarised towards C2 (resp., C1), where C2 denotes the anti-vaccine community, and C1 the pro-vaccine community. Source: Schmidt et al. (2018).

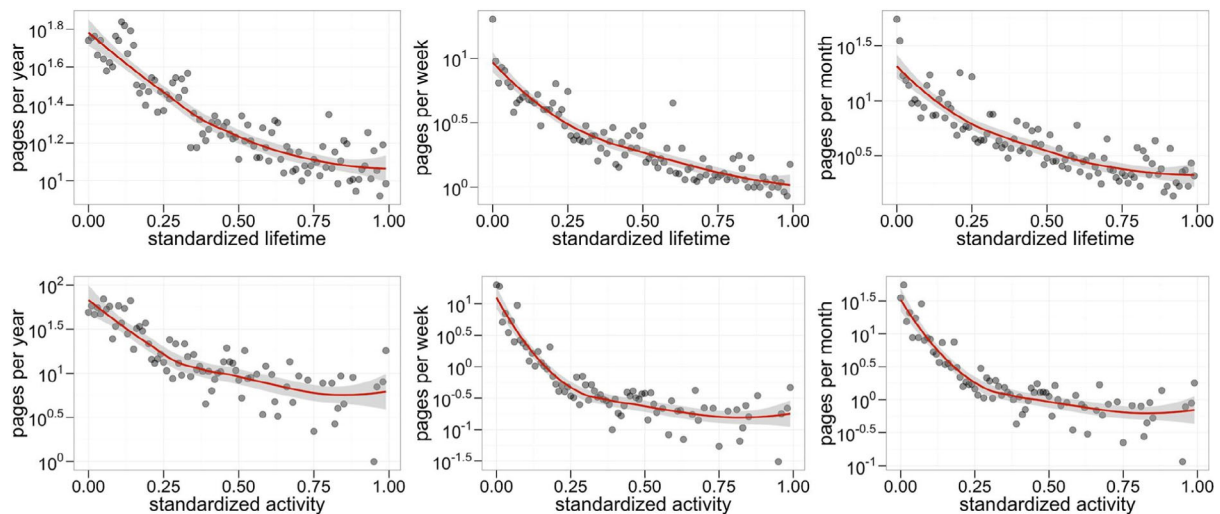
Figure 7: Vaccines. Probability density function (PDF) of users' polarisation $\rho(u)$ considering likes

3.6. News consumption and selective exposure

So far, we have measured the echo chambers effect by analysing information with respect to specific narratives, such as science versus pseudoscience, Brexit or vaccines. What we have observed so far is that users tend to look for information that already fulfils their system of beliefs and form polarised groups around a shared world view. In this section, we generalise our previous results by focusing on how users interact with information produced by official news providers on Facebook, by analysing the news consumption patterns of 376 million users over a time span of 6 years (Schmidt et al., 2017).

On the basis of previous results, it should be clear that it is the very nature of the users – possibly fostered by Facebook News Feed algorithm – to promote the emergence of echo chambers. Indeed, here again, we observed the exact same dynamics, finding that users are strongly polarised, and that

their attention was confined within specific clusters (communities) of pages. To quantify the turnover of news sources, we measured the heterogeneity of users' activity, i.e. the total number of pages with which a user interacts. Figure 8 shows the number of news sources that a user interacts with during her lifetime – i.e. the temporal distance between the first and last interactions with a post – and for increasing levels of engagement (total number of likes). For a comparative analysis, both lifetime and engagement were standardised between 0 and 1. Results are shown for the yearly time window (first column), for the weekly (second column), and the monthly (third column). We may notice that a user usually interacts with a small number of news outlets and that higher levels of activity and longer lifetime corresponded to a smaller number of sources. Indeed, while users with very low lifetime and activity levels interacted with about 100 pages in a year, 30 in a month and 10 in a week, the same values are far lower for more active and long-lived users, who interacted with about 10 pages in a year, and less than four monthly and weekly. Users naturally tended to confine their activity to a limited set of pages, so proving that news consumption on Facebook is dominated by selective exposure.



Conclusions: Highest lifetime and activity are denoted by value 1, lowest by 0. Source: Schmidt et al. (2017).

Figure 8: Maximum number of unique news sources with which users with increasing levels of lifetime (top) and activity (bottom) interacted (yearly, weekly and monthly, respectively)

In this manuscript, we present a collection of works investigating how information is consumed and shapes communities on Facebook. We found that confirmation bias plays a crucial role in content selection and diffusion, showing that users are highly focused on their specific narrative. Moreover, we provide empirical evidence of the existence of echo chambers, i.e. well separated and polarised groups of like-minded users sharing the same narrative. Immersed in these bubbles, users keep framing and reinforcing their world view, and may end up in a more extreme position (Sunstein, 2002). In this scenario, corrections in the form of fact-checking or debunking attempts seem to fail and have instead a backfire effect, reinforcing users' pre-existing beliefs.

Empirical evidence suggests that misinformation spreading on social media is directly related to the increasing segregation of users in echo chambers. To contrast misinformation, smoothing polarisation is therefore essential. To this end, users' behaviour may be used as a proxy to determine in advance the targets for hoaxes and fake news (Del Vicario et al., 2019). A timely identification of potential misinformation targets may allow the design of tailored counter-narratives and appropriate communication strategies. In this direction, we are now working to analyse, design, test and evaluate different strategies to improve science communication on social media, with special attention to delicate and polarising topics that need to be addressed with care, such as vaccines or climate change.⁷

To conclude, it is also worth noting that the same social networks – grounded on the paradigm of like – might feed and foster echo chambers, and therefore users' polarisation. For the first time in history, the stock market is dominated by five technology companies, i.e. Apple, Alphabet, Microsoft,

⁷ Ongoing H2020 QUEST project: <https://cordis.europa.eu/project/rcn/219700/factsheet/en>

Amazon and Facebook (Barwise and Watkins, 2018). As of December 2018, Facebook has 2.32 billion monthly active users,⁸ and is now planning to integrate WhatsApp, Instagram and Facebook Messenger, bringing together more than 2.6 billion users.⁹ Important questions arise about platforms' regulation, as well as concerns about data privacy and security, control and the potential manipulation of information. Therefore there is the need for increasing platforms' transparency and data accessibility. Currently, it is very difficult or even impossible, to access social media data for research purposes, and this is a major barrier to quantitative research and analysis (Taylor et al., 2018). In this respect, the collaboration of social media platforms is vital and highly encouraged.

References

- Adamic LA and Glance N, 2005. The political blogosphere and the 2004 US election: divided they blog. In: Proceedings of the 3rd International Workshop on Link Discovery. ACM, pp. 36–43.
- AGCOM, 2018. News vs. Fake in the Information System. Interim Report Sector Inquiry "Online platforms and the news system".
- Barwise P and Watkins L, 2018. The Evolution of Digital Dominance. How and Why We Got to GAFA. In: Digital dominance: the power of Google, Amazon, Facebook and Apple. Oxford University Press, Oxford, UK.
- Bessi A, Zollo F, Del Vicario M, Scala A, Caldarelli G and Quattrociocchi W, 2015. Trend of narratives in the age of misinformation. PLoS ONE, 10, e0134641.
- Bessi A, Petroni F, Del Vicario M, Zollo F, Anagnostopoulos A, Scala A, Caldarelli G and Quattrociocchi W, 2016. Homophily and polarization in the age of misinformation. The European Physical Journal Special Topics, 225, 2047–2059.
- Bhagat S, Burke M, Diuk C, Filiz IO and Edunov S, 2016. Three and a half degrees of separation. In Facebook Research Blog: <https://research.fb.com/three-and-a-half-degrees-of-separation/>
- Coe K, Kenski K and Rains SA, 2014. Online and uncivil? Patterns and determinants of incivility in newspaper website comments. Journal of Communication, 64, 658–679.
- Del Vicario M, Bessi A, Zollo F, Petroni F, Scala A, Caldarelli G, Stanley HE and Quattrociocchi W, 2016. The spreading of misinformation online. Proceedings of the National Academy of Sciences, 113, 554–559.
- Del Vicario M, Zollo F, Caldarelli G, Scala A and Quattrociocchi W, 2017. Mapping social dynamics on Facebook: the Brexit debate. Social Networks, 50, 6–16.
- Del Vicario M, Quattrociocchi W, Scala A and Zollo F, 2019. Polarization and fake news: Early warning of Potential misinformation targets. ACM Transactions on the Web (TWEB), 13. <https://doi.org/10.1145/3316809>
- Ellison NB, Steinfield C and Lampe C, 2007. The benefits of Facebook "Friends": social capital and college students' use of online social network sites. Journal of Computer-Mediated Communication, 12, 1143–1168.
- Howell WL, 2013. Digital Wildfires in a Hyperconnected World. Global Risks Tech. Report. World Economic Forum.
- Lévy P and Bononno R, 1997. Collective intelligence: mankind's emerging world in cyberspace. Perseus Books, New York, USA.
- Nic N, Fletcher R, Kalogeropoulos A, Levy DA and Nielsen RK, 2018. Reuters Institute Digital News Report 2018. Reuters Institute for the Study of Journalism.
- Nyhan B and Reifler J, 2010. When corrections fail: the persistence of political misperceptions. Political Behavior, 32, 303–330.
- O'Hara K and Stevens D, 2015. Echo chambers and online radicalism: assessing the internet's complicity in violent extremism. Policy & Internet, 7, 401–422.
- Schmidt AL, Zollo F, Del Vicario M, Bessi A, Scala A, Caldarelli G, Stanley HE and Quattrociocchi W, 2017. Anatomy of News Consumption on Facebook. Proceedings of the National Academy of Sciences, 114, 3035–3039.
- Schmidt AL, Zollo F, Scala A, Betsch C and Quattrociocchi W, 2018. Polarization of the vaccination debate on Facebook. Vaccine, 36, 3606–3612.
- Serrano MA, Boguna M and Vespignani A, 2009. Extracting the multiscale backbone of complex weighted networks. Proceedings of the National Academy of Sciences, 106, 6483–6488.
- Steinberger R, Pouliquen B and Van der Goot E, 2013. An Introduction to the Europe Media Monitor Family of Applications. arXiv preprint arXiv:1309.5290.
- Sunstein CR, 2002. The law of group polarization. Journal of Political Philosophy, 10, 175–195.
- Sunstein CR, 2012. Breaking up the echo. NY Times, A25.
- Taylor S, Pickering B, Grace P, Boniface M, Bakir V, Boyd D, Engesser S, Epstein R, Fawzi N, Fernbach P, Fisher DR, Gardner BC, Jacobs K, Jacobson S, Krämer B, Kucharski A, McStay A, Mercier H, Metzger M, Polletta F, Quattrociocchi W, Sloman S, Sperber D, Spierings CHBM, Wardle C, Zollo F and Zubiaga A, 2018. Opinion forming in the digital age. <https://doi.org/10.5281/zenodo.1468575>
- Travers J and Milgram S, 1977. An experimental study of the small world problem. Sociometry, 32, 425–443.

⁸ <https://zephoria.com/top-15-valuable-facebook-statistics/>

⁹ <https://www.nytimes.com/2019/01/25/technology/facebook-instagram-whatsapp-messenger.html>

- Zollo F and Quattrociocchi W, 2018a. Misinformation spreading on Facebook. In: *Complex Spreading Phenomena in Social Systems*. Springer, Cham. pp. 177–196.
- Zollo F and Quattrociocchi W, 2018b. *Social Dynamics in the Age of Credulity: the misinformation risk and its fallout*. In: *Digital Dominance: The Power of Google, Amazon, Facebook and Apple*. Oxford University Press, Oxford, UK.
- Zollo F, Novak PK, Del Vicario M, Bessi A, Mozetič I, Scala A, Caldarelli G and Quattrociocchi W, 2015. Emotional dynamics in the age of misinformation. *PLoS ONE*, 10, e0138740.
- Zollo F, Bessi A, Del Vicario M, Scala A, Caldarelli G, Shekhtman L, Havlin S and Quattrociocchi W, 2017. Debunking in a world of tribes. *PLoS ONE*, 12, e0181821.

Abbreviations

- EMM Europe Media Monitor
PDF probability density function
WEF World Economic Forum