



<https://helda.helsinki.fi>

Helda

The Spread of Anti-vaccination Memes on Facebook

Knuutila, A

2024

Knuutila, A, George, A, Bright, J & Howard, P 2024, The Spread of Anti-vaccination Memes on Facebook. in M Preuss, A Leszkiewicz, JC Boucher, O Fridman & L Stampe (eds), Disinformation in Open Online Media : 6th Multidisciplinary International Symposium, MISDOOM 2024, Münster, Germany, September 2–4, 2024, Proceedings. Springer, Cham, pp. 86-100, Disinformation in Open Online Media. 6th Multidisciplinary International Symposium, MISDOOM 2024, Münster, Germany, 02/09/2024. https://doi.org/10.1007/978-3-031-71210-4_6

<http://hdl.handle.net/10138/587590>

10.1007/978-3-031-71210-4_6

cc_by

publishedVersion

Downloaded from Helda, University of Helsinki institutional repository.

This is an electronic reprint of the original article.

This reprint may differ from the original in pagination and typographic detail.

Please cite the original version.



The Spread of Anti-vaccination Memes on Facebook

Aleksi Knuutila¹ (✉), Anna George², Jonathan Bright³, Anna George²,
and Philip Howard²

¹ University of Helsinki, Unioninkatu 35, 00170 Helsinki, Finland
aleksi.knuutila@helsinki.fi

² Oxford Internet Institute, 1 St Giles', Oxford OX1 3JS, UK

³ Turing Institute, 96 Euston Road, London NW1 2DB, UK

Abstract. False claims about vaccines can find large audiences online, leading to vaccine hesitancy. The most influential content on social media is often visual, but studies about misinformation largely focus on text instead of images. This study uses new image analysis capabilities that Facebook and Instagram have made available to understand the spread of visual anti-vaccination memes on these platforms. We identified 200 influential memes that contain scepticism or hesitancy towards vaccines and the 15,000 public Facebook accounts on which the memes have been shared. We describe the memes' spread on a large scale by identifying communities of accounts and describing the diffusion pathways of memes between the communities. We develop a novel method of testing whether a meme has spread from one community of accounts to another that works on sequential time series alone. We identify 16 distinct communities of Facebook accounts and categories them based on thematic and regional focus. Anti-vaccination memes originate predominantly from North American Facebook accounts. These accounts often focus on opposing COVID-19 policies or promoting conspiracy theories about elites. Memes from these communities also spread internationally, particularly to Europe, demonstrating their influence beyond North America. The analysis demonstrates that memes receive the most engagement within their initial community. However, their overall reach depends on their ability to spread to other communities. This suggests that the ability of memes to find large audiences is based on their capacity to spread beyond their original contexts and to be used by groups with potentially different agendas.

Keywords: Anti-vaccine activism · memes · Facebook · network structures · information diffusion

1 Introduction

Over the past decade, anti-vaccination rhetoric has become part of the mainstream discourse regarding the public health practice of childhood vaccination [6]. Its proponents utilise social media to foster online spaces that strengthen and popularise anti-vaccination discourses [12]. One influential form of media in these communities is memes - static

images, typically with some text layered on top, which can be used to communicate simple, persuasive concepts. Memes can have a significant impact because they often gather large audiences, communicate simple, powerful messages and elicit strong emotional reactions. They are simple visual forms of communication which can potentially “go viral” and spread beyond the original communities of interest they were created in.

Describing the spread of memes is particularly difficult as they are often copied from one place and posted to another without any attribution, obscuring their origin and the steps they take moving from one place to another. Hence, it has been difficult to scientifically study memes’ origin and propagation patterns on a large scale. Visual material on the Internet in general been studied much less than textual forms of information [11]. In contemporary debates about vaccinations, when new types of claims and evidence are emerging all the time, visual and memetic forms of communication deserve particular attention.

We use time series data about posts containing anti-vaccination memes to answer the following questions:

- What communities do anti-vaccination memes transmit through?
- How important is the crossing from one community to another for the visibility of anti-vaccination memes?

One challenge in studying memes is that there is often a lack of meta-data about the source of the content. In our paper, we use timestamp information and time series analyses to detect when the meme first appeared on the platform and how it spread through the platform. This approach allows us to map the large-scale patterns of diffusion of anti-vaccination memes, from American conservative accounts to European and African groups, in a statistically robust and interpretable form.

2 Related Work

Visual media have distinct advantages over text. Images and image-text combinations offer what Shifman calls “high information density” [16], i.e. messages in this form can be understood at a glance. Images can elicit strong emotional reactions in people, especially if they display something graphic, such as the use of violence. Memes have frequently been used to refer to particular kinds of images, though the concept more broadly encompasses all cultural artefacts. In an influential article on the concept, Wiggins and Bowers offer a broader definition of a meme, seeing it as “remixed or parodied spreadable media” [20]. The term “spreadable media” points to the capacity of contemporary digital media to spread quickly and gather large audiences when large numbers of people share it on their networks. Spreadable media would include, for instance, YouTube videos, which can be viewed millions or even billions of times through the sharing of links. The type of meme shared on social media at a given time can correspond to life events or current internet culture ‘fads’. What is specific to memes is the remixes or alterations that the media undergo when they spread. It can be difficult to find the “original authors” of memes as each person can alter the text or image of a meme and then spread their version to their network. However, the various versions of memes can still convey similar information or be related to a similar topic. This article focuses on visual still images (e.g., memes) with anti-vaccination content.

Social media exhibits the phenomenon of homophily, i.e. the fact that relationships form more likely between people who are similar in some sense [14]. An obvious example of this is that information tends to spread between accounts that share a common language or common interests. As a result, the network of accounts has a community structure, i.e. there are clusters of accounts that are more densely connected between themselves and have fewer connections with other accounts. This structure, in turn, affects the way information, including memes, travels through the network, for instance, by limiting the global flow of information since leaps between the communities are less likely [15].

Social media data has made it possible to model the effects of community structure on the flow of memes. Weng et al. [19], for instance, demonstrate that jumping between many communities early in a meme’s life cycle predicts its future popularity. Several studies have also looked at the structure of specifically vaccination-related communities online. Studies on vaccination information overall on Twitter often find a single anti-vaccination community with little contact with pro-vaccination users [11]. Johnson et al. [9] find many distinct communities of anti-vaccination users on Facebook. They claim that the anti-vaccination communities are localised geographically, for instance, based on US states, but share information effectively through additional communities with a global orientation.

Memes have often been described as originating from relatively small communities of dedicated hobbyists before jumping to other demographics and large audiences. De Zeeuw et al. [4], for instance, discuss a diffusion pattern they name “normification”, in which a cultural artefact from a fringe online subculture finds a larger and more dispersed audience, for instance, through going viral on Facebook or being replicated in newspapers or broadcast media. In this work, they found that platforms such as YouTube and Reddit were “bridge platforms”, meaning they brought fringe ideas to the mainstream, including mainstream news media [4].

3 Methodology

3.1 Data Collection

This article uses data from CrowdTangle [2]. CrowdTangle is a service operated by Facebook that tracks Facebook, Instagram, and Reddit posts. This report uses data about text in images to search for memes. Since 2018, Facebook has done optical character recognition on a large scale on images shared on both Facebook and Instagram [17]. The platform uses optical character recognition to determine what text images contain and publishes this information as part of the image’s metadata. CrowdTangle collects this information in its database and makes it possible to search for images based on the text contained therein. This makes it possible to find all posts containing particular memes in CrowdTangle’s database by searching for distinctive phrases that are part of the memes. The CrowdTangle database contains public accounts on Instagram, pages, groups, and verified profiles on Facebook. For the sake of simplicity, we refer to pages, groups and profiles as “Facebook accounts” in this article. CrowdTangle’s database is not comprehensive, but it contains almost all public Facebook pages and groups with more than 50,000 followers and a significant share of smaller accounts [3].

In the first step of the research, we identified 200 memes with anti-vaccination themes with a relatively large audience. We searched for images on Crowdtangle based on the text content in the images and limited the search to one year, from the 15th of February 2020 to the 15th of February 2021. Notably, this period is mostly focused on the time before actual vaccination programmes, which in Europe and the United States began in December 2020. The results, however, point to the fact that there was an active discussion of the vaccines even prior to the vaccination programmes. The keywords we used are “vaccination”, “vaccinate”, “vaccine”, “pfizer”, and “astrazeneca”. For each day this year, we obtained 2,000 images that contained text that matched the keywords and had received an exceptional amount of engagement given the account they had been posted in (i.e., were “overperforming” in CrowdTangle’s terminology). The results from these searches were aggregated by the image texts, i.e. we identified the text contents of images that occurred frequently in this dataset.



Fig. 1. An example of variations of an anti-vaccination meme in which the text content stays the same, but the visual template changes

We identify anti-vaccination memes from this dataset. We reviewed the most frequently reposted images and selected anti-vaccination memes with the following criteria:

- The images express opposition or scepticism towards vaccines, including but not limited to vaccines against the coronavirus.
- We include statements that vaccines are not necessary for the lack of harm from pathogens, the description of vaccination side effects, expressing hesitancy because of questionable motivations attributed to the government or pharmaceutical companies, or general negative statements related to the vaccine (such as “Down with vaccines”).
- We also include humorous or satirical statements in this category when they imply a stance against vaccinations, even when this stance is not intended sincerely. This humour could take, for instance, the form of mocking public figures associated with vaccines, such as Fauci or Gates, for their work on the vaccine.
- We exclude images that are screenshots of social media posts or screenshots of news articles where these do not have any user-generated alterations made to them.

One research team member went through images in the Crowdtangle search results until 200 distinct anti-vaccination memes had been identified. This number of distinct memes was identified after examining 2065 images. The coding of this researcher was verified through double coding a sample of 100 decisions, with 50 of both images that had been coded as anti-vaccination memes and images that had not. The Cronbach's alpha indicator describing the reliability of the coding was 0.81, indicating good reliability in the coding.

For each anti-vaccination meme we identified, we made global searches on Crowdtangle based on the meme's text content. This allowed us to find all the copies of the same images shared on Facebook and variations of the same meme where the visual content had changed, but the text had remained the same. Figure 1 provides an example of variations in a meme where the text stays the same. In this example, all three images would have been grouped as instances of the same meme. Memes with the same image but different text would have been grouped as distinct memes.

3.2 Hierarchical Clustering of Accounts

The 200 anti-vaccination memes in our dataset have been shared in 14,594 Facebook accounts. Our analysis focuses not on individual accounts on Facebook but rather on what we call communities, i.e. groups of accounts that are densely connected in terms of frequently sharing posts from one account to another. The focus on communities will make large-scale diffusion patterns interpretable and provide insights that would be lost by focusing on, for instance, the individual most influential accounts.

Our community detection method is based on forming a graph in which the nodes are Facebook accounts, and the edges are formed based on posts being shared from one account to another. To create this graph, we do not use posts exclusively related to anti-vaccination memes but all posts shared between the accounts. This is because few memes shared between accounts have information about their origins, and using the full set of posts shared between accounts reflects the affinity between accounts better.

Some Facebook posts contain information about where the contents of the posts are shared from. The source of the post's content is recorded when people share it from one account to another using Facebook's internal sharing functionality (similar to retweets). Notably, this information about the source of shared information is not always available. For instance, for our dataset of anti-vaccination memes, information about the source of the posts was available 29% of the time. This metadata is insufficient for unbiased, reliable inference about the transmission of anti-vaccination memes specifically. However, the information about the source of posts is enough to understand the larger overall pattern of connections between accounts, at least when there are enough data points available.

We download up to 5,000 posts from each account from the previous 12 months to create the graph. From these posts, we identify shares of posts that had originated from other Facebook accounts. Based on such shares, we form a directed graph with weights reflecting the frequency of shares. The resulting graph contains 147,188 edges, an average of 10.0 edges for every account. We want to focus our analysis on accounts with a substantial number of connections to other accounts by taking the 10th k-core of this graph, i.e. focusing on the maximal subgraph that contains nodes with ten edges

or more. This subgraph has 5,405 accounts with 115841 edges. We selected the 10th k-core because we estimate that for accounts with less than ten connections to other accounts, we cannot reliably estimate which community individual accounts should be connected to accounts with less than ten connections to other accounts. We chose ten as a relatively high and conservative threshold, given that there is some randomness and uncertainty about which post shares on Facebook contain the metadata about the source of the content.

We use Ward's method [13] to identify a hierarchy of potential communities. Ward's method is agglomerative, i.e. it starts from all observations forming separate communities and chooses which communities to merge at each step following an objective function. In our study, we create node embeddings based on generalised singular value decomposition embeddings at ten dimensions. We then apply Ward's minimum variance criterion as an objective function, i.e. merging communities at each step that increases the within-communities variance the least. The resulting hierarchy describes many different ways the data can be clustered and how individual communities can be broken up or combined depending on the desired level of resolution.

4 Identifying Transmission Patterns Without Metadata About Sources

As mentioned above, for a large share, the metadata about the source of their content is unavailable. To study the transmission of anti-vaccination memes from one community to another, we develop a method that does not rely on metadata about sources but works based on time series data describing when specific memes appeared in specific communities. We use time series data that describes the first appearance of every 200 memes within every community (if they appeared in the community at all). We estimate the spreading pattern by looking at whether memes appear systematically in one community before appearing in another. This pattern, a meme in one community systematically preceding another, does not prove an actual causal connection, as confounding factors might cause it. Nonetheless, it is a pattern that we would necessarily observe if memes spread from one community to another.

Our estimate of the spreading pattern is based on comparing observations to a null model. For the null model, we assume that the frequency of each meme posted to each community is the same as that observed, but the sequence of the postings is random. With statistical tests, we can derive the likelihood that the meme appears in one community before another for any meme and pair of communities.

In the null model, the likelihood that the meme will first appear in a particular community is:

$$p_{X,m} = \frac{N_{X,m}}{N_{all,m}}, \text{ where } N \text{ is the number of memes } m \text{ in the community } X$$

The likelihood of a meme appearing in one community X it appears in community Y in the null model can be simplified to

$$p_{X,Y,m} = \frac{p_{X,m}}{p_{X,m} + p_{Y,m}}$$

The equations give us the expected frequency at which the meme appears in one community before another if the posts in one community do not influence other communities at all. We assess the likelihood that the distribution of the expected and observed frequencies is the same using the G-test or the log-likelihood ratio test [21]. This statistical test is performed for every pair of two distinct communities in a set of communities. This means that the number of statistical tests is high. To correct the so-called multiple testing problem, we apply the Benjamini-Hochberg procedure to adjust p-values [18]. The subsequent analysis discusses adjusted p-values.

4.1 Semi-supervised Selection of Communities

Above, we described the hierarchical community detection method, which results in a hierarchy that describes how individual communities can be combined or split apart. Based on the hierarchy, it is possible to choose many different sets of communities for analysis. For our analysis, we developed a novel approach for selecting sets of communities, applying the method of identifying transmission to the community hierarchy. We call this approach semi-supervised because it combines the results from the unsupervised hierarchical community detection we performed with Ward’s method with the additional prior knowledge about potential communities, specifically the time series data about the appearance of anti-vaccination memes.

We search for the best selection of a set of communities, using the communities hierarchy from Ward’s method as a set of options. Our evaluation of a particular selection of communities is based on the certainty with which we can trace the movement of memes from one community to another. We want to select sets of communities with explanatory power for describing the transmission of memes, i.e. sets of communities for which there is statistically significant transmission according to the methodology described above. Moreover, we aim to select sets of communities where as many as possible have both incoming and outgoing connections of transmission. This is done to identify rich and descriptive networks between communities. We hence divide every community’s connections to other communities into incoming and outgoing connections. This difference is based on comparing observed data with the null model described above. Where a meme in a particular community precedes a meme appearing in another community at a frequency higher than expected by the null model, we classify the connection as an outgoing connection and vice versa.

We formalise this selection process with a cost function, to reward the selection of communities with both incoming and outgoing connections that are statistically significant. The cost function takes a sum of all p-values above the significance threshold and adds a constant alpha value for every p-value below the threshold. Alpha is a parameter that defines how much weight to give to significant connections relative to reducing the p-value of non-significant connections. For this work, we used an alpha value of -0.5.

We use the following cost function:

$$Cost = |S| \cdot \alpha + \sum_{p \in NS} p$$

where:

$$S = p \in P(I) \cup P(O) \mid p < \text{threshold}$$

and

$$NS = p \in P(I) \cup P(O) \mid p \geq \text{threshold}$$

Here, (I) and (O) represent the set of all incoming and outgoing connections where observed < expected and observed > expected, respectively. (P(I)) Moreover, (P(O)) represent the lowest p-value of these connections. (α) is the alpha value, (|S|) is the number of significant p-values.

Given this cost function, we search through the hierarchy of communities and identify the set of communities for which the cost function gives the lowest value. We start from the top of the hierarchy (with two communities containing all the accounts in the dataset) and recursively evaluate whether the selection is improved by replacing individual communities with their subcommunities in the community hierarchy. We run the beam search tree traversal algorithm to go through the options in the community hierarchy recursively [10]. For our study, we used a beam width of 3 and traversed 50,000 options for sets of communities.

The Girvan-Newman modularity score attained with Ward’s method was 0.64. However, the resulting set of communities only has three statistically significant connections between the communities. This means that an unsupervised community detection method does not result in a choice of communities that can be used to explain the spreading of memes. In contrast, the semi-supervised method finds 19 statistically significant connections between the communities. The resulting set of communities has a Girvan-Newman modularity score of 0.54.

5 Results

5.1 Describing the Resulting Communities

The methodology we described resulted in 16 distinct communities, encompassing 5405 Facebook accounts. To describe the communities, we read through the list of 20 accounts that had shared the largest amount of anti-vaccination memes in each community. Additionally, we examined the most distinctive words in the names of accounts in each community according to the metric of Scaled F-Scores [7]. Two researchers worked on the interpretation to increase the validity of the community description.

The results are summarised in Table 1. For readability, we first categorised the regional focus of every community. The largest share of the communities were focused on the United States or Canada. Some of the communities also contained largely English-language accounts with an international focus, while some also had a range of African, Australian or European accounts. Moreover, we categorised the communities based on their thematic content. Many of the communities contained predominantly accounts focused on communicating and coordinating opposition to COVID-19 restrictions (“Re-Open California #EndTheLockdown”) or discussing what we called elite conspiracies (e.g. “Exposing the Satanic World Government”). Many North American communities focused on supporting individual conservative politicians or parties, most prominently Donald Trump. Especially in Europe, we found many communities containing accounts of political extra-parliamentary protest movements, such as the Yellow Vests.

Table 1. 16 communities of Facebook accounts, with a classification of their regional focus and thematic focus, distinctive words in account names and descriptive statistics

Label	Region	Type	Distinctive uni/bi-grams in account names	# accounts	# posts with memes
A1	US/Canadian	Elite conspiracy		8	65
A2	US/Canadian	Anti-Covid policy	trudeau, canadian, alberta, masks, freedom fighters	261	1011
A3	US/Canadian	Anti-Covid policy	wrong, awaken, cosmic, save america, natural	178	594
A4	US/Canadian	Conservative politics	lake, unfiltered, angry, fake, politically incorrect	214	505
B1	International	Elite conspiracy	flat earth, exposing the, david icke, collective, elite	610	2309
B2	US/Canadian	Conservative politics	vote, trump train, deplorable, americans for, joe biden	921	2137
B3	US/Canadian	Conservative politics	wisconsin, excessive quarantine, against excessive, reopen, protect	122	300
B4	US/Canadian	Conservative politics	president trump, america great, protest, nevada, utah	479	1329
C1	International	Elite conspiracy	black lives, jamaica, positive, hip hop, conscious	396	882
C2	International	Elite conspiracy		37	126
C3	African	Cultural	anc, eff, zimbabwe, marcos, duterte	381	776
C4	African	Cultural	nigeria, zambia, biafra, malawi, university	741	1152

(continued)

Table 1. (continued)

Label	Region	Type	Distinctive uni/bi-grams in account names	# accounts	# posts with memes
C5	Australian	Conservative politics	aussie, australia, stand up for, stop the, sovereign	148	479
C6	American	Conservative politics	ted cruz, michigan, latinos, indiana, liberty memes	284	630
C7	European	Anti-systemic	gilets jaunes, vegan, romania, france, italia	396	702
C8	European	Anti-systemic	ireland, nederland, waarheid, tegen, yellow vest	229	556

We described such groups based on the concept of anti-systemic movements [8]. In particular, African communities contained a relatively wide variety of accounts that were not overtly political but focused on questions such as religion, parenting and nutrition. We categorised these accounts under the category “Cultural”.




Fig. 2. A network graph displaying the detected movement of memes between communities of accounts, with the thematic and regional classification of the communities

5.2 The Spreading Pattern of Anti-vaccination Memes

Figure 2 shows a graph of connections between the communities, visualised using the Dagre layout [5]. We draw edges in the graph in all cases where there is a statistically significant connection between the communities, i.e. the test described above suggests that memes spread from one community to another. The significance test is based on aggregate data about all of the 200 anti-vaccination memes we identified. In this graph, the communities in the topmost row only have outgoing connections, the middle row has outgoing and incoming connections, and the bottom row has only incoming connections. We also labelled the communities to reflect their position in this network. Communities whose label starts with A are upstream in the network, i.e., they have outgoing edges online. In contrast, communities with the label B have both in-going and outgoing edges, and communities with the label C have only in-coming edges (i.e. they are downstream in the movement of memes).

Our results indicate that anti-vaccination memes originate from communities with predominantly North American accounts. The communities upstream from other communities include those focused on Anti-Covid policies (A2 and A3) and a community focused on elite conspiracies (A1). One upstream community (A4) was focused on conservative politics, but to a large extent, the conservative politics communities were further downstream. This suggests that the online spaces on Facebook where memes are first shared or where they are created are within accounts focused on opposing COVID-19 policies or those focused on a range of elite-related conspiracy theories. The presence of the A1 community in the graph suggests that very small communities (8 Facebook accounts in the case of A1) can have a large influence since this community has statistically significant connections to three other communities focused on elite conspiracies and two communities focused on conservative politics.

Figure 2 also describes a clear regional pattern. The memes in our study originate exclusively from groups associated with the United States and Canada. We identify some international, African, Australian and European communities that are further downstream. Connections with Europe and Africa illustrate that accounts that focus on American conservative politics and opposition to COVID-19 policies have influence and find an audience also outside of the region of North America. It is also notable that European communities have statistically significant connections with several different thematic types of communities. This suggests that there is no single pathway but multiple connections through which anti-vaccination memes spread to accounts in Europe.

5.3 The Impact of Spreading Between Communities

In this section, we describe how important it is for memes to cross boundaries for their lifecycle and popularity. In our dataset, memes traverse an average of 10.6 of the 16 communities, with a standard deviation of 3.3. Since our sampling strategy emphasised trending memes, many of them can be expected to move successfully between communities. It takes an average of 5.7 days for a meme to leave the first community in which it appears.

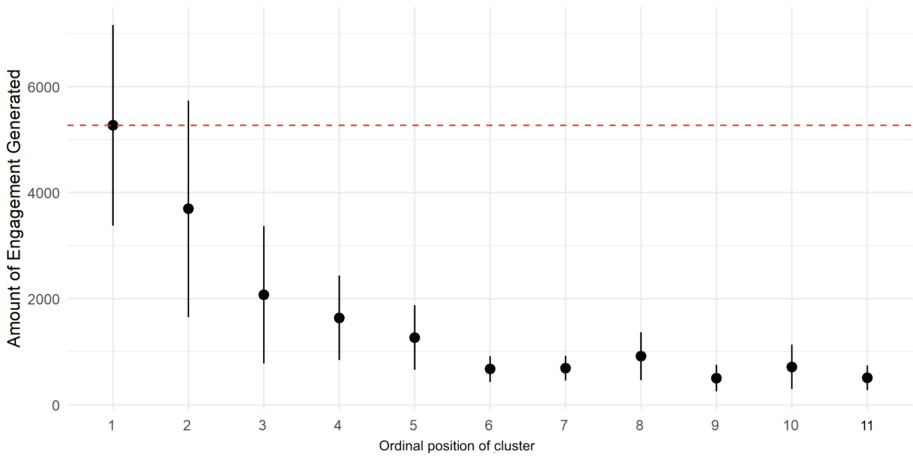


Fig. 3. The amount of engagement received by anti-vaccination memes in the first and subsequent communities in which they appeared

Figure 3 describes the amount of engagement (i.e. reactions such as likes or shares) generated by the first communities in which memes appear and by every subsequent community. Since Facebook does not provide metrics such as view counts that would directly describe the amount of attention posts have received, we use engagement as a proxy for the visibility of posts. Figure 3 shows the first community where the meme accounts for a significant share of their total engagement. The memes also receive engagement from other communities but with diminishing returns. Differences in the relative sizes of communities do not explain the effect. While the mean number of accounts in the upstream communities (A1–A4) is 165.3, for downstream accounts (C1–C8) it is 494.7.

Even though memes receive the most engagement from the first cluster in which they appear, their overall reach still depends on them crossing clusters. Figure 4 describes how much of the meme’s reach can be attributed to the first of the clusters. It shows that 32.1% of all engagement (across all clusters) results from posts in the first cluster. For the metric tracking the total number of posts that contained the meme, the share of the first cluster was smaller, 24.0%. These figures show that the potential reach of the memes is significantly higher if they can cross onto other clusters. The metric called duration measures the difference in time between the meme’s first and last appearance in a particular cluster. It shows that memes survive and get reposted relatively long in their first cluster, but not as long as all clusters combined. The duration of memes in the first cluster was 59.1% of their duration overall.




Fig. 4. The share of total engagement, audience, duration and number of posts for anti-vaccination memes that is attributed to the first community in which they appear

6 Discussion and Limitations

The results show that anti-vaccination memes typically first appeared on Facebook on accounts that opposed COVID-19 policies or focused on conspiracy theories about elites. From there, memes spread towards accounts associated with Conservative politics, especially the Republican party in the US and conservative parties in Canada and Australia. The results also show a regional pattern. Particularly US and Canadian accounts have international influence, as their memes spread into Europe, Africa and Australia. This paper contributes to the literature on online anti-vaccination discourse and the study of memes by being among the first to describe such large-scale regional patterns. We also describe a novel methodology for studying the spread of memes based on time-series data and the memes' text content.

There appears to be a change in the types of groups they reach when crossing regions. Particularly in Europe, we found they were associated with anti-systemic, extra-parliamentary political movements, while in Africa, the memes were shared in a wide range of “cultural” accounts. The difference between Europe, the US, and Canada may reflect that there typically was a broader consensus between parliamentary parties concerning lockdown and vaccination policies during the pandemic [1].

The results speak to the power of memes as a medium to cross between contexts and groups with very different agendas. In this respect, memes are like what Tuters calls “floating signifiers”; they are single images that can be used for various purposes and are “nevertheless temporarily united through affective bonds.” [19]. This ability to move between different audiences is also what makes memes so influential. Our findings support this conclusion. They show that the first community in which memes appeared only accounted for about a third of the overall engagement they received.

The study has some limitations that future studies could overcome. We group memes based on their text content, mainly because this is a technical possibility offered by our data source, CrowdTangle. The use of CrowdTangle also meant some biases in the sample, especially the exclusion of smaller Facebook pages and groups. In some

memes, the visual template stays the same while the text content changes. A different way of grouping memes may give different results. The study only covers 200 memes, which we identified by looking at trending posts based on English-language keywords. The keywords could have included other languages or the trademarks of vaccines that are specific to some regions. The identified patterns could have been different if we included a more extensive set of memes or other languages. We also describe different communities of accounts relatively coarsely through five categories. Another paper could explore in more detail the variety of accounts contained in every community and describe comparatively the differences between communities of accounts. In this study, we also did not have the opportunity to study the actual content of memes and how the contents may influence how they spread on social networks.

Notes on Funding. Kate Joynes-Burgess is funded by a Wellcome Trust PhD Studentship in the Humanities and Social Sciences.

References

1. Bol, D., Giani, M., Blais, A., Loewen, P.J.: The effect of COVID-19 lockdowns on political support: some good news for democracy? *Eur. J. Polit. Res.* **60**, 497–505 (2021). <https://doi.org/10.1111/1475-6765.12401>
2. CrowdTangle Team. CrowdTangle. Facebook, Menlo Park, California, United States. List ID: 1432547 (2021)
3. CrowdTangle: What data is CrowdTangle tracking? (2021). <http://help.crowdtangle.com/en/articles/1140930-what-data-is-crowdtangle-tracking>
4. De Zeeuw, D., Hagen, S., Peeters, S., Jokubauskaite, E.: Tracing normification. *First Monday* **25** (2020). <https://doi.org/10.5210/fm.v25i11.10643>
5. Gansner, E.R., Koutsofios, E., North, S.C., Vo, K.-P.: A technique for drawing directed graphs. *IEEE Trans. Softw. Eng.* **19**, 214–230 (1993). <https://doi.org/10.1109/32.221135>
6. Garrett, R., Young, S.D.: Online misinformation and vaccine hesitancy. *Transl. Behav. Med.* **11**, 2194–2199 (2021). <https://doi.org/10.1093/tbm/ibab128>
7. Gutierrez-Bustamante, M., Espinosa-Leal, L.: Natural language processing methods for scoring sustainability reports—a study of Nordic listed companies. *Sustainability* **14**, 9165 (2022). <https://doi.org/10.3390/su14159165>
8. Holt, K.: Alternative media and the notion of anti-systemness: towards an analytical framework. *Media Commun.* 49–57 (2018). <https://doi.org/10.17645/mac.v6i4.1467>
9. Johnson, N.F., et al.: The online competition between pro- and anti-vaccination views. *Nature* **582**, 230–233 (2020). <https://doi.org/10.1038/s41586-020-2281-1>
10. Lemons, S., Linares López, C., Holte, R.C., Ruml, W.: Beam search: faster and monotonic. *ICAPS* **32**, 222–230 (2022). <https://doi.org/10.1609/icaps.v32i1.19805>
11. Marchal, N., Neudert, L.-M., Kollanyi, B., Howard, P.N.: Investigating visual content shared over Twitter during the 2019 EU parliamentary election campaign. *Media Commun.* **9**, 158–170 (2021). <https://doi.org/10.17645/mac.v9i1.3421>
12. Milani, E., Weitkamp, E., Webb, P.: The visual vaccine debate on Twitter: a social network analysis. *Media Commun.* **8**, 364–375 (2020). <https://doi.org/10.17645/mac.v8i2.2847>
13. Murtagh, F., Legendre, P.: Ward’s hierarchical agglomerative clustering method: which algorithms implement Ward’s criterion? *J. Classif.* **31**, 274–295 (2014). <https://doi.org/10.1007/s00357-014-9161-z>

14. Newman, M.: *Networks: An Introduction*. Oxford University Press, Oxford
15. Onnela, J.-P., et al.: Structure and tie strengths in mobile communication networks. *Proc. Natl. Acad. Sci.* **104**, 7332–7336 (2007). <https://doi.org/10.1073/pnas.0610245104>
16. Shifman, L.: *Memes in Digital Culture*. MIT Press (2013)
17. Sivakumar, V., Gordo, A., Paluri, M.: Rosetta: understanding text in images and videos with machine learning (2018). <https://engineering.fb.com/2018/09/11/ai-research/rosetta-understanding-text-in-images-and-videos-with-machine-learning/>
18. Thissen, D., Steinberg, L., Kuang, D.: Quick and easy implementation of the Benjamini-Hochberg procedure for controlling the false positive rate in multiple comparisons. *J. Educ. Behav. Statist.* **27**, 77–83 (2002). <https://doi.org/10.3102/10769986027001077>
19. Weng, L., Menczer, F., Ahn, Y.-Y.: Virality prediction and community structure in social networks. *Sci. Rep.* **3**, 2522 (2013). <https://doi.org/10.1038/srep02522>
20. Wiggins, B.E., Bowers, G.B.: Memes as genre: a structural analysis of the memescape. *New Media Soc.* **17**, 1886–1906 (2015). <https://doi.org/10.1177/1461444814535194>
21. Woolf, B.: The log likelihood ratio test (the G-test). *Ann. Hum. Genet.* **21**, 397–409 (1957). <https://doi.org/10.1111/j.1469-1809.1972.tb00293.x>

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

