



Misinformation and Polarization around COVID-19 vaccines in France, Germany, and Italy

Gianluca Nogara
gianluca.nogara@supsi.ch
University of Applied Sciences and
Arts of Southern Switzerland
Switzerland

Francesco Pierri
francesco.pierri@polimi.it
Politecnico di Milano
Italy

Stefano Cresci
stefano.cresci@iit.cnr.it
IIT-CNR
Italy

Luca Luceri
lluceri@isi.edu
USC Information Sciences Institute
USA

Silvia Giordano
silvia.giordano@supsi.ch
University of Applied Sciences and
Arts of Southern Switzerland
Switzerland

ABSTRACT

The kick-off of vaccination campaigns in Europe, starting in late December 2020, has been followed by the online spread of controversies and conspiracies surrounding vaccine validity and efficacy. We study Twitter discussions in three major European languages (Italian, German, and French) during the vaccination campaign. Moving beyond content analysis to explore the structural aspects of online discussions, our investigation includes an analysis of polarization and the potential formation of echo chambers, revealing nuanced behavioral and topical differences in user interactions across the analyzed countries. Notably, we identify strong anti- and pro-vaccine factions exhibiting heterogeneous temporal polarization patterns in different countries. Through a detailed examination of news-sharing sources, we uncover the widespread use of other media platforms like Telegram and YouTube for disseminating low-credibility information, indicating a concerning trend of diminishing news credibility over time. Our findings on Twitter discussions during the COVID-19 vaccination campaign in major European languages expose nuanced behavioral distinctions, revealing the profound impact of polarization and the emergence of distinct anti-vaccine and pro-vaccine advocates over time.

KEYWORDS

COVID-19, vaccine, disinformation, misinformation, social media, echo-chamber

ACM Reference Format:

Gianluca Nogara, Francesco Pierri, Stefano Cresci, Luca Luceri, and Silvia Giordano. 2024. Misinformation and Polarization around COVID-19 vaccines in France, Germany, and Italy. In *ACM Web Science Conference (WEBSCI '24)*, May 21–24, 2024, Stuttgart, Germany. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3614419.3644020>



This work is licensed under a Creative Commons Attribution International 4.0 License.

WEBSCI '24, May 21–24, 2024, Stuttgart, Germany
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0334-8/24/05
<https://doi.org/10.1145/3614419.3644020>

1 INTRODUCTION

Digital platforms played a central role in the dissemination of news and information during the COVID-19 pandemic [10]. While online social networks such as Twitter (currently X) provide open and immediate forums for discussion, they also expose their users to abusive behavior and harmful content [13, 34]. An infodemic of false and misleading content [4, 5] has likely had a negative impact on vaccination campaigns around the world, contributing to an increase in the number of people who are skeptical and in denial about COVID-19 and the related vaccine [27, 39].

The viral spread of online misinformation regarding COVID-19 vaccines posed a serious threat to the safety of many people, as digital platforms are often designed to maximize user engagement [28] and often serve as an important vector in the spread of fake news and conspiracy theories [2, 35, 43]. Social media platforms also contribute to the creation of user groups that are bound together and hostile toward users with opposing views and opinions [44], leading to increased polarization among users and the creation of large anti-vaccine groups [20]. During the COVID-19 pandemic, European countries acted together in order to swiftly develop safe and effective vaccines, with the first COVID-19 vaccination programs kicking off in late December 2020, less than a year into the crisis [41].

In this work, we explore the landscape evolution of online vaccine discussions in three European languages: French, German, and Italian. In particular, we address the following research questions (RQs):

- RQ1:** *Do we observe the presence and formation of echo chambers? Is there an increase in network polarization over time?*
- RQ2:** *What is the prevalence of low-credibility information in online discussions around vaccines? What are the main vectors for the proliferation of unreliable content?*

To answer these questions, we draw on a large-scale dataset of over 30M tweets related to COVID-19 vaccines collected over a period of one year. We study behavioral differences and similarities in the vaccine discourse across countries by conducting a study of user interactions in a dynamic (i.e., time-varying) fashion, providing the following contributions:

- We studied the formation of groups of anti- and pro-vaccination supporters by conducting a temporal study of the interaction networks, showing that opposing communities tend to become progressively polarised over time.
- We analyzed the prevalence of news from low-credibility sources through the use of fact-checking and media reliability services, uncovering a downward trend in information quality in all considered countries over time.
- We carried out a cross-media analysis of some of the main social media used to disseminate fake news [1, 26], characterizing their content and contextualizing their use.

2 RELATED WORK

2.1 COVID-19 misinformation

The vaccine campaign against COVID-19 that kicked off in Europe in December 2020 is situated in a historical period marked by the spread of misinformation, defined as ‘an era of fake news’ [46]. It is, therefore, important to shed light on how users on social media have handled the discussion. Research shows how a handful of users can flood social media platforms with questionable content and receive massive endorsement [35, 45]. A case in point is represented by the Disinformation Dozen, responsible for spreading fake news related to COVID-19 and its vaccine around the world [35]. These actions have not only the ultimate effect of manipulating public opinion on social media but can also lead to offline repercussions for people’s health, e.g., creating fear and distrust in vaccine administration [37], or decreasing vaccination coverage and increasing the health risk of exposed people [27].

The debate related to vaccination has already shown that it is an argument subject to the formation of opposing user factions. Cossard et al. [6] showed how echo chambers are created in Italy in the debate related to measles based on pro-vax and no-vax ideology. Another study looking at Italy [7], this time studying the debate related to the anti-COVID-19 vaccine, showed the division there is between supporters and individuals hesitant to continue during the vaccination campaign, yet found a possible common set of facts on which the two sides could agree. Other studies, such as [25], consider multilingual data and show how the COVID-19 pandemic has led the no-vax communities to become more central to country-specific debates, forming a global Twitter network against vaccination. Political components in discourse have also been highlighted in the debate, [19] conducted a study in COVID-19 conversations in the United States, showing that they are largely influenced by political bias.

2.2 Vectors of Misinformation: Echo Chambers and Multimedia Platforms

Meanwhile, there is also a growing interest in studying how online user interactions evolve over time, since such evolution is linked to the dynamics of interactions and to the formation of echo chambers [40]. Among the studies on the structure of networks and their communities, [30] suggests that anti-vaxxer communities tend to be denser and more organized than pro-vaxxer communities.

Moving beyond a content-based analysis, [29] studied more than just the content of the data but also investigated the structure of the network and its characteristics, identifying four different and

unconnected discussion clusters of potentially influential individuals capable of manipulating public opinion, showing the ease with which they generate content and become influential. Others focused on temporal analyses, such as [16] that analyzed discussions about the COVID-19 vaccine in the US, studying the variation of vaccination topics on Twitter from December 2020 to February 2021, showing not only a significant temporal variation but also a different attitude among different communities.

Other studies have focused on the potential vectors that have been used for sharing false or misleading news. Among them, [26] found that over a quarter of the most viewed YouTube videos on COVID-19 contained misleading information, reaching millions of viewers worldwide and creating downstream risks for viewers who use YouTube as a source of information. [11] suggests that this trend is not only related to the COVID-19 pandemic, but it is a growing phenomenon already noted in the past. Indeed, the study draws attention to the fact that most of the videos were negative in tone and that the annual number of uploaded videos has increased. Another platform known to hold anti-vaxxer content is Telegram, which presents itself as an unmoderated platform [1]. Focusing on users, far-right groups and COVID-19 protest groups emerged [8], interlacing far-right discourse in COVID-19 protest groups.

Given these premises, our objective is to conduct an in-depth examination of the phenomenon of misinformation and user biases within a dataset centered on discussions about COVID-19 vaccines in Europe. This study affords us the opportunity to explore the structure and external media employed in online conversations surrounding COVID-19 vaccines on social media. An initial investigation was carried out on the existing data, involving only descriptive and statistical analyses [9], including hashtags, news sources, and geolocation, serving as an initial foundation for further studies akin to the one presented here.

3 DATA

We rely on a multilingual dataset of vaccine-related conversations on Twitter, totaling more than 30 million tweets collected using the Twitter Historical API from November 2020 to June 2021 [9], a period which covers the start of vaccination campaigns in European countries. The dataset consists of online discussions around the COVID-19 vaccine on Twitter in three major European languages (French, German, and Italian).¹ For the sake of simplicity, in our analyses, we will refer to the three European countries in which these languages are spoken the most, namely France, Germany, and Italy, although the data might capture conversations from other countries as well (e.g., Austria, Canada, Switzerland, etc.). Figure 1 shows the distribution of different tweets shared by users in our dataset, with the large majority composed of retweets in all three languages.

As an exploratory step, we look at the prevalence of different hashtags, which provides insights into the discussions that have engaged users in each country. Figure 2 shows the top 10 most used hashtags, highlighting how the main topics of discussion were those related to vaccines and COVID-19. However, for all three datasets, we also find hashtags related to the political sphere. For example, we have *#Macron* for the French language dataset,

¹For more details on the data collection, we refer the reader to the related paper [9].




Figure 1: Distribution of different kinds of tweets shared for French data (FR), German data (DE), and Italian data (IT).

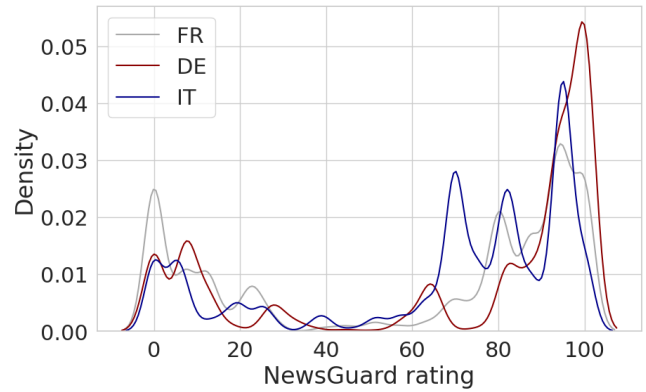

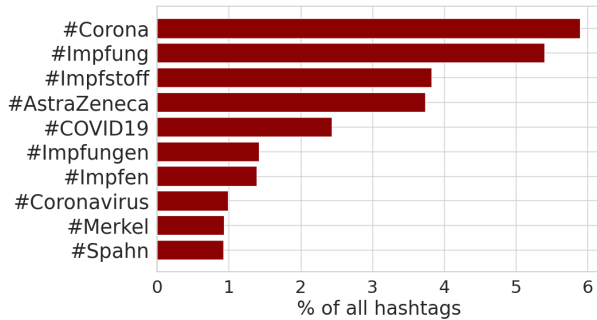



Figure 3: Density distribution of reliability ratings of tweets containing a domain classified by NewsGuard, estimated by Kernel Density Estimation (KDE), for each country.



(a) French dataset



(b) German dataset



(c) Italian dataset

Figure 2: Top-10 most frequent hashtags by country.

#Draghi for the Italian language dataset, #Merkel and #Spahn for the German language dataset, highlighting how despite the differences between countries we find rather common behavioral patterns.

Also of importance is #AstraZeneca, which stands out as a highly used hashtag in all three datasets, remarking how much impact AstraZeneca’s discontinuation of the anti-COVID-19 vaccine has had in public opinion [18].

4 METHODOLOGY

4.1 Reliability of Information Sources

To ensure a comprehensive and reliable assessment of the credibility of news outlets shared along vaccine-related content, we used a distant-supervision approach to label websites adopted in the literature [36, 37, 45]. For the labels, we employed NewsGuard² ratings, given its established reputation as an independent and transparent organization that employs human experts in the field of journalism. NewsGuard conducts rigorous evaluations of news sources, considering factors such as transparency, accountability, adherence to journalistic standards and error correction, and promoting high-quality information [38]. We used the service via an API that provides credibility ratings for news outlets, with scores ranging from 0 to 100, where 0 indicates a completely unreliable news outlet, and 100 is a completely trustworthy one. Specifically, we extracted all URLs from the tweets in our dataset and we queried the API to obtain credibility scores for more than 50% of all URLs mentioned in our datasets, retrieving a score for 65% of all URLs in the French dataset, 70% of all URLs in the German dataset, and 57% of all URLs in the Italian dataset. Since NewsGuard provides various reliability ranges [32], we classify domains into two categories: high credibility and low credibility. To select a threshold to determine whether or not a news outlet is credible, we analyzed the distributions of ratings across our datasets. As shown in Figure 3, the distributions of NewsGuard credibility ratings for the three datasets tend to be bimodal. In fact, we find a large number of domains whose credibility ratings are ≥ 60 and ≤ 30 . Following [37], we used a conservative criterion and labeled those news outlets with NewsGuard rating ≤ 30 as *low credibility*, and those with rating ≥ 60 as *high credibility*.

²<https://www.newsguardtech.com/>

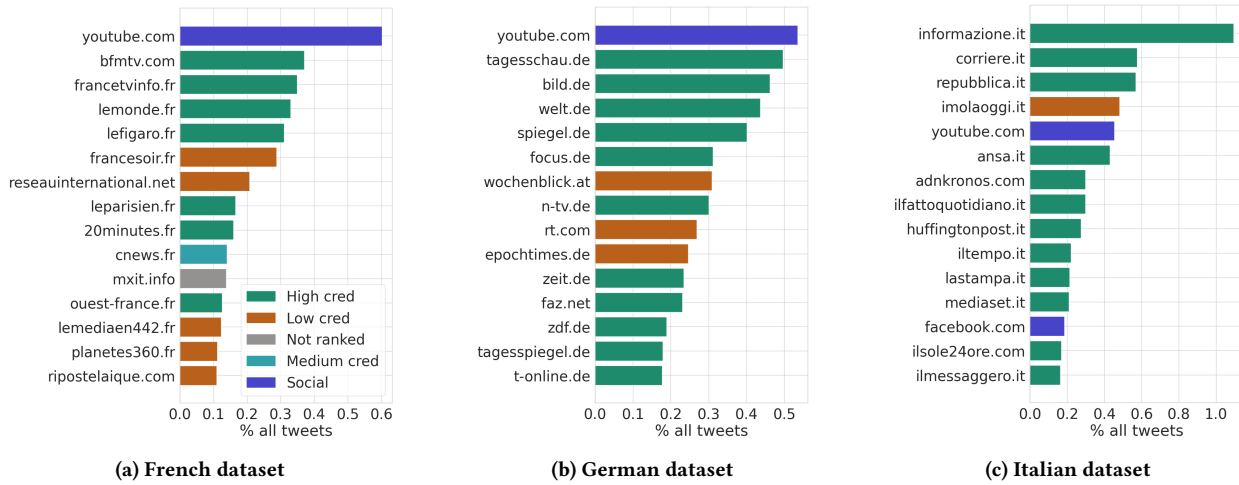


Figure 4: Top-15 shared domains in each dataset colored according to their NewsGuard rating.

An exploratory analysis of most shared web domains is provided in Figure 4. We observe a large prevalence of YouTube, which is particularly relevant given the role of this platform for spreading fake news and conspiracy theories [15, 17, 24]. We also notice the presence of at least one low-credibility web domain in the top-15 ranking of all three countries.

Given the importance of YouTube links in our datasets, which account for 2.2% (171.256) of the links shared, we exploited the YouTube Data API to collect information on linked videos. We were able to obtain a range of information such as the availability of a video, as well as its duration, title, description, number of views, number of likes, number of comments, channel, publication date, and the possible presence of a label related to COVID-19 [22, 23].

4.2 Building Retweet Networks

To analyze the structure of interactions between users in our datasets, we built and analyzed a retweet network for every country. As retweets often indicate a form of social endorsement [31], they represent an effective method to identify groups of users interacting with and approving each other [40]. To study how these users are connected and grouped within the retweet network, we apply the well-known Louvain community discovery algorithm [3]. Since we aim to study how the interactions evolve over time, we built and compared three different snapshots of the interaction networks, each taken at a different point in time. These times were chosen to capture peaks of activity corresponding to important real-world events. We consider a 7-day lag from the event for gathering the bulk of tweet cascades about the events, which we use to build our retweet networks.

Figure 5 shows the selected intervals with respect to Twitter activity in the different datasets. We can see similar temporal patterns across languages, with peaks corresponding to significant events such as the start of COVID-19 vaccine administration in Europe on December 27, 2020, commonly referred to as the “Vaccine Day”, and the suspension of the AstraZeneca vaccine by several European



Figure 5: Moving average of the daily number of tweets shared in each dataset. We highlight periods in which we constructed three different retweeting networks with different colors.

countries (e.g., Italy, France, Germany, Spain, Portugal and Austria) on March 15, 2021. As these events are not only national but also European in character, the patterns in the number of activities across the three datasets are quite consistent. There is an increase in the number of interactions during international events, and minor variations in activity levels can be attributed to secondary or internal factors.

We considered the following intervals to build three different data snapshots, for each language:

- T1:** [01/11/2020 - 11/01/2021] a week after the peak of activity on January 5th, which corresponds to the beginning of the vaccination campaign;
- T2:** [12/01/2021 - 31/03/2021] a week after the peak of activity on March 15th, which corresponds to the suspension of AstraZeneca in many European countries, including France, Germany and Italy;
- T3:** [01/04/2021 - 30/06/2021] end of the collection period.




Figure 6: Number of tweets captured in each snapshot for French (FR), German (DE), and Italian datasets (IT).

Each snapshot contains a substantial amount of user data and tweets in order to construct and analyze retweet networks. Figures 6 and 7 show the numbers of tweets and users underlying each snapshot. Networks corresponding to different periods might be overlapping in terms of nodes, in case users are active during multiple periods.

5 RESULTS

5.1 RQ1: Polarization

By building a retweet network for each dataset at different time intervals T1, T2, and T3, we study the evolution of user interactions and the possible formation of communities, or the confluence of smaller communities in those of greater relevance [35, 40]. We analyze and classify clusters of users based on their shared content and interactions. By manually examining the hashtags and texts used by these users, we categorize them within specific clusters. Our study is not solely focused on nodes with the highest centrality or popularity; rather, we randomly select nodes to investigate and categorize their shared content. Analyzing the content shared by the different network communities of the Italian language dataset, in fact, we find that the anti-vaxxer community contains very few political figures. This finding relates to the Italian government’s situation during most Italian vaccination campaigns. In fact, the election of a new prime minister (i.e., Mario Draghi), who remained in office from February 13, 2021 to October 22, 2022 led to a government of national unity without real opposition. The situation of the French dataset is significantly different given the French political situation during the period under analysis. The government led by Emmanuel Macron received strong criticism, especially from far-right and left-wing parties, such as *Les Patriotes*, whose leader belongs to the anti-vaxxer community in the French network.

In addition to analyzing the content of the communities, we also applied metrics to define the degree of fragmentation and overall polarization of the networks at each time interval. In literature, Random Walk Controversy (RWC) is a well-known measure that assesses the degree of fragmentation and controversy within polarized discussion networks [21]. We then computed RWC scores at each snapshot and for each network, as reported in our annotations to Figure 8.

In Figure 8, we highlight the evolution of the retweet networks at times T1, T2, and T3. At each time step, each node is colored




Figure 7: Number of users captured in each snapshot for French (FR), German (DE), and Italian datasets (IT).

based on the community it belongs to, according to the following scheme:

- the blue color represents the community characterized by a strong *pro-vaxxer* component;
- the orange color represents the community characterized by a strong *anti-vaxxer* component;
- the grey color represents *other* communities.

Through a qualitative analysis of user-shared content and user clusters, we were able to characterize the datasets. By performing a manual study of the textual content shared by the most active users and random users within the network. We then labeled the users and clusters according to their feeds, content, and the users that comprise them, so we were able to identify the peculiar dynamics of online debates about vaccination as they unfolded over time in the three countries, which can be summarized as follows:

- The French language dataset is characterized by several political communities, with pro-vaxxer users coinciding with the ruling parties and media, and anti-vaxxer users with government opposition groups. This dataset is thus characterized by a strong political component.
- The German language dataset is very homogeneous and generally does not have well-defined user communities, which become evident during the last period of analysis (T3).
- The Italian language dataset maintains the same structure from the first to the last snapshot, with two large communities of anti-vaxxer and pro-vaxxer users mainly clashing for ideological reasons rather than for political ones.

Except for the German network, anti-vaxxer and pro-vaxxer communities appear early in the discussion, as shown by their presence at T1 for Italian and French. However, the three networks exhibit a similar structure at T3, when about 80% of all users either belong to the anti-vaxxer or pro-vaxxer communities. This kind of behavior leads communities to close outward and generate echo chambers in which tweets are shared by users with similar interests, values, or identities who then tend to aggregate [20, 42].

We can observe that the RWC scores for French and Italian exhibit similar high values throughout the period of analysis (~0.8). The RWC score for German increases from 0.204 to 0.938, showing the formation of two echo chambers, one anti-vaxxer and one pro-vaxxer, which are only evident in the last period of analysis (T3).

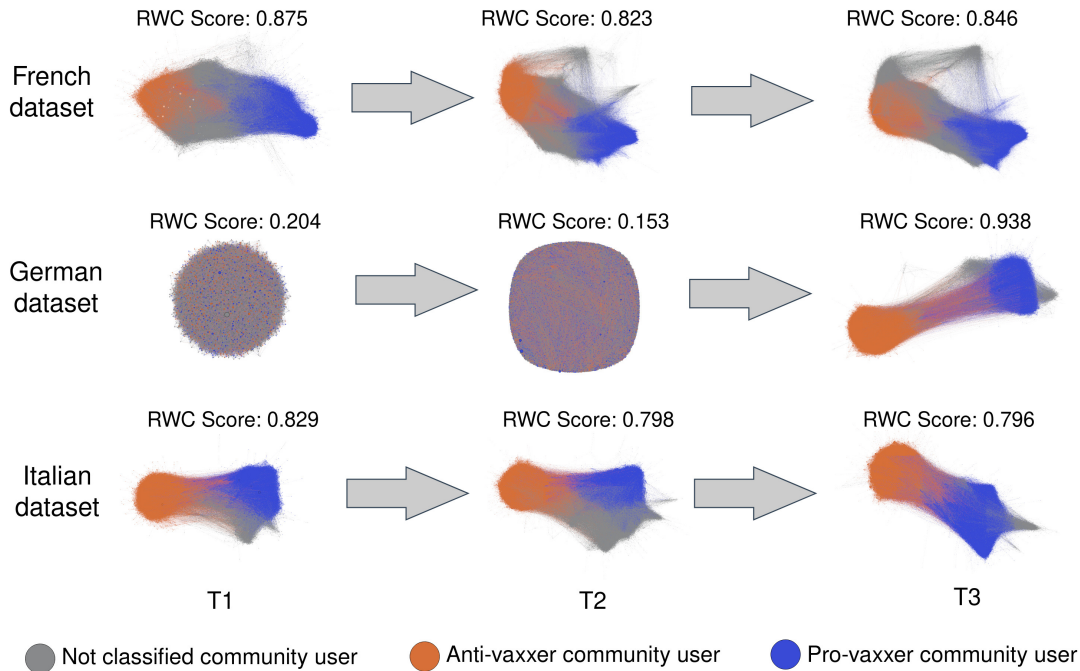


Figure 8: Temporal evolution of retweet networks at different time intervals with related controversy scores (Random Walk Controversy – RWC) for each country.

5.2 RQ2: Media Credibility

Links to external websites are widely shared by users in our dataset, with more than 6,500,000 tweets containing at least one URL (excluding retweets that link to "twitter.com" or ".co"). We analyze these URLs to better understand the type of content that is shared by users, both in terms of the type of link shared (e.g., to a website or another social network) and the reliability of the external website.

5.2.1 NewsGuard and source reliability. We first analyze the average reliability score of news websites shared by Twitter users over time. As shown in Figure 9, we observe that in the three datasets the average daily consumption of news leans toward reputable sources (FR = 61.0, DE = 70.5, IT = 67.2). We remark that according to NewsGuard, a reputable source has a score higher than 60. We can notice that the reliability of news websites in French conversations about vaccines is smaller compared to Italian and German ones.

Next, we analyze the daily prevalence of tweets sharing links to either High- or Low-credibility news websites in the three datasets. Figure 10 shows that tweets sharing news articles are a minority (less than 10% of daily tweets on average), but reputable sources are significantly more prevalent than unreliable ones in all datasets (cf. median values in the caption). In line with previous findings, we notice that low-credibility information makes up a small but not negligible proportion of online sources [37].

We then focus on the reliability of news articles shared by anti-vaxxer and pro-vaxxer communities. Anti-vaxxer communities are often responsible for spreading disinformation and misleading news [14] about vaccines.




Figure 9: Average daily rating of links to websites on NewsGuard's list.

As we can see from Figure 11, pro-vaxxers groups, in all countries, mainly share links referring to high-credibility sources. The links referring to low-credibility sources, both in the French and in the Italian language dataset, is a very minimal percentage, close to zero. And this is valid for all time periods. Only in the German dataset we observe an increase in the error bar due to the more evolving nature of this dataset. The situation in anti-vaxxers communities see a very strong number of links referring to low-credibility sources. Interestingly, the anti-vaxxers communities identified in the French language dataset show a prevalence of links referring to low-credibility sources, suggesting a strong presence of disinformation. In all three dataset, we see a high variability along time (large error bar) as we can observe an increasing polarization in the anti-vax discussion in all languages.




Figure 10: Distributions of the daily proportion of tweets linking to High- and Low-credibility news websites for each dataset. Median HIGH NewsGuard rating FR 5.5, DE 7.5 and IT 7.3, median LOW NewsGuard rating FR 2.7, DE 2.6 and IT 1.8

Table 1: Stats related to the NewsGuard rating of domains of users who share YouTube links .

Dataset	Rating (all users)	Rating (users sharing YT)
FR	62.0	46.4
DE	70.5	56.1
IT	67.2	53.6

5.2.2 Links to Other Media Platforms. We investigate whether YouTube plays a role in promoting false and misleading claims about vaccines, as often reported in literature [17, 26, 35], by extracting metadata about videos shared by Twitter users in the three datasets.

We first look at the prevalence of COVID-19 labels attached to videos by the platform as well as their availability at the time of the analysis (May 2023). COVID-19 labels has been placed on all videos that have COVID-19 as a topic of discussion in 2021. This label leads back to the WHO website. The purpose of this action is to counter misinformation on the platform [12]. The unavailability of a video refers to a suspension of the channel, a suspension of the video, or the author deleting the video from the platform. Unavailable videos have been linked to sources of misinformation in the past [37, 45]. As shown in Figure 12, over 60% are related to COVID-19 vaccines, as expected, given the topic of conversations. Approximately 50% of the videos in French and Italian are no longer available, while almost 70% are still available in German. This suggests a higher prevalence of unreliable content originating on YouTube in French and Italian than in German.

Through a qualitative study, we identified several YouTube channels related to anti-vax movements among those most shared by users in all countries. To achieve this, all links related to YouTube videos were filtered from the data by going to the most re-shared videos and studying the channels that posted them in order to get a clearer view of the content and type of channels. We then investigate the reliability of news sources shared by users who shared YouTube videos, finding that such users tend to share less reliable

news articles, as shown in Table 1. To corroborate this analysis, we perform a co-occurrence analysis on the bipartite network of users and their shared URLs. A projection of this bipartite graph produces a co-occurrence weighted network whose weight between URLs i and j represents the total number of users that shared both i and j . Next, we assign the NewsGuard score of the website to each URL, if available. For all YouTube URLs, we can thus extract a reliability score by taking the weighted average score of all its neighbors. Figure 13 shows the distribution of such YouTube reliability scores for all datasets. We can notice that the median values are very low in all cases (FR = 20, DE = 12.5, and IT= 20.0), further indicating that YouTube URLs are often shared by users actively engaging with unreliable news sources [24].

A similar analysis was carried out on the shared Telegram channels, which might be very relevant vectors of misinformation since they are not moderated in any way [1, 33]. Telegram links, which account for 0.16% (12.395) of the links shared, allowed us to get the information and content shared by different public channels. We obtained 492 channels and 4,714,193 messages (62 channels with 628,548 messages in the French dataset, 324 channels with 2,607,947 messages in the German dataset, and 106 channels with 1,477,698 messages in the Italian dataset). From the point of view of shared messages, the channels in the German-language dataset have the highest volume; however, the channels in the Italian-language dataset have the highest ratio of message per channel. A manual qualitative inspection indicates that most shared channels from all datasets are mainly from anti-vaxer or conspiracy groups, involving more than one million users. Through NewsGuard, we obtain the rating of domains shared within Telegram channels, similar to the rating of links shared by the users in the Twitter datasets.

Figure 14 shows that links sharing news articles within Telegram channels tend to outnumber unreliable ones. The distance between the medians is, however, smaller than the data observed in Figure 10.

6 DISCUSSION

6.1 Conclusions and future works

We carried out a detailed study of Twitter conversations in French, German, and Italian related to the COVID-19 vaccination campaigns between 2020 and 2021. We investigated the trustworthiness of news domains shared in this period, identifying a low but non-negligible prevalence of tweets sharing unreliable sources compared to mainstream and reputable sources (which are shared 2-4x more times). We also studied the formation of pro-vaxer and anti-vaxer user groups with their respective echo-chambers over time.

The findings from RQ1 indicate that the German language dataset exhibits a distinct pattern compared to other datasets, with a notable absence of significant polarization and a more homogeneous discussion regarding COVID-19, which polarizes only during the end of the period of analysis. To confirm this hypothesis, we have shown how there are substantial differences in the polarization of users in the different datasets, shown again how the German language dataset tends to polarize after a time interval compared to the other datasets. However, we can take note that although there are differences in the discussion given by political and behavioural aspects, we have shown how all datasets tend to converge into a




Figure 11: Percentage of tweets linking to High- and Low-credibility news websites for the *anti-vaxxers* and *pro-vaxxers* communities obtained from the retweet network for each dataset. We consider data from all iterations, and the error bar is computed with the discrepancy in each time period T

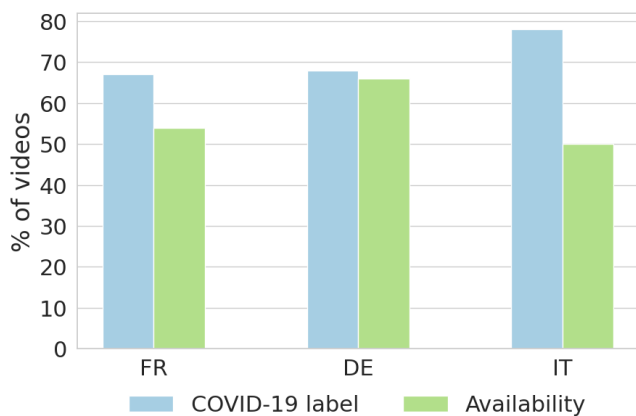


Figure 12: Proportion of YouTube videos shared on Twitter that had a COVID-19 label (blue) and those that are still available (green) as of May 2023, for each dataset.

well-defined structure despite completely different starting points and communities.

The construction of retweet networks allows us to generalize the dynamics of interaction and information dissemination on platforms such as Twitter. This approach allows us to analyze patterns of connection between users, trends in content dissemination, and the evolution of conversations over time.

Analysis of retweet networks reveals how, despite the different cultural origins and specific issues addressed in the various datasets, entirely similar structures emerge at the last snapshot. These converging patterns suggest that there are universal mechanisms that drive the dissemination of information and the formation of communities within Twitter, although discussions may differ depending on context.




Figure 13: Distributions of reliability ratings (NewsGuard) of domains closely related to YouTube, as derived from the weighted network analysis, for each dataset.

Through the analysis of retweet networks, we are able to gain a comprehensive and deep view of the social and informational dynamics on Twitter, transcending language and cultural barriers to identify the trends and behaviors that characterize this global communication platform on a given topic.

Regarding the prevalence of low-credibility information and how they are circulated (RQ2), we explored the credibility of the links to external websites shared in the tweets. By using NewsGuard to score the credibility of the linked domains, we observed again an evolution: The percentage of low-credibility domains shared within the anti-vaxxer communities is already dominant, like in the French dataset, or increases with time to prevail, like in the Italian and German discussion. For the pro-vaxxer communities, we found an high prevalence of high-credibility links, stable in the French and Italian datasets or with a tendency to improve in the German discussion. Not surprisingly, we also observed a growth of the links to external social media, like YouTube and Telegram that are a clear indication of increasing disinformation, as a symptom of the rise of polarization.

In the future, we also aim to extend our comprehensive analyses




Figure 14: Distributions of the percentage of links per Telegram channel linking to high and low credibility news websites for each dataset. Median HIGH NewsGuard rating FR 9.7, DE 11.8 and IT 8.4, median LOW NewsGuard rating FR 5.1, DE 5.5 and IT 4.6

by considering extending the temporal scope of the study to more recent data would provide a deeper understanding of the evolution of vaccine discourse over time. It would also be useful in addition to the three major EU languages, examining vaccine discourse in other languages and regions could highlight broader cross-cultural patterns. This could help identify patterns that we can call global and generalize user behaviors in a given context such as vaccines.

6.2 Limitations

Our work presents some limitations. First, the set of keywords used for data collection could introduce a degree of bias, favoring certain perspectives over others related to the vaccines. Additionally, our choice to obtain the data using only keywords might result in an incomplete representation of the topic, omitting crucial nuances and emerging trends that might not be covered by the selected terms as well as linguistic and cultural contexts that employ different terminologies.

We rely on data collected through Twitter’s Historical Search API V2 for all three datasets in order to ensure overall consistency in the collected data. However, it is important to note that this choice involves a trade-off. While it provides us with stability and consistency in the data over time, it results in the exclusion of suspended users or removed content that would have been obtained through the use of the Filter API in real-time streaming mode. As a result, the possible absence of such removed or suspended content emerges as a limitation of our investigation.

An additional limitation of our study arises from the use of keywords in three specific languages: Italian, French, and German. Although this has enabled us to gather data across a wide range of countries and regions where these languages are spoken, it is essential to recognize that this methodology might capture not only content from the native countries of these languages but also from nations where they are used as secondary or intercultural communication languages. A clear example is provided by the dataset in the French language, which exhibits a significantly higher number

of activities compared to the others. This data is undoubtedly influenced by the presence of numerous French-speaking countries worldwide. As a result, the content collected could reflect both the realities of the languages’ countries of origin and those of the regions where they are employed.

6.3 Ethical Considerations

We did not attempt to identify or deanonymize individual users, and we do not share any personal information and political leanings reporting only public content and figures. The data presented was collected through public APIs in an aggregate manner, and the ID of tweets is provided transparently so that access can be granted in full accordance with the platforms’ data-sharing policies; except for posts that have been removed or made private by users, which hinders reproducibility analyses. At the time of this writing, we acknowledge that Twitter has limited access to its public API, and we encourage interested researchers to contact us should they be interested in working with our datasets.

REFERENCES

- [1] Aleksi Knuutila Aliksandr Herasimenka, Jonathan Bright and Philip N. Howard. 2023. Misinformation and professional news on largely unmoderated platforms: The case of Telegram. *Journal of Information Technology & Politics* 20, 2 (2023).
- [2] Edmund Andrews. 2019. How fake news spreads like a real virus. <https://engineering.stanford.edu/magazine/article/how-fake-news-spreads-real-virus>
- [3] Vincent D Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefevre. 2008. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment* 2008, 10 (2008).
- [4] Andrea Calamusa, Serena Tardelli, Marco Avvenuti, Stefano Cresci, Ilaria Federigi, Maurizio Tesconi, Marco Verani, and Annalaura Carducci. 2020. Twitter monitoring evidence of Covid-19 infodemic in Italy. *European Journal of Public Health* 30 (2020).
- [5] Matteo Cinelli, Walter Quattrociocchi, Alessandro Galeazzi, Carlo Michele Valensise, Emanuele Brugnoti, Ana Lucia Schmidt, Paola Zola, Fabiana Zollo, and Antonio Scala. 2020. The COVID-19 social media infodemic. *Scientific Reports* 10, 1 (2020).
- [6] Alessandro Cossard, Gianmarco De Francisci Morales, Kyriaki Kalimeri, Yelena Mejova, Daniela Paolotti, and Michele Starnini. 2020. Falling into the Echo Chamber: The Italian vaccination debate on Twitter. *Proceedings of the International AAAI Conference on Web and Social Media (ICWSM)* (2020).
- [7] Giuseppe Crupi, Yelena Mejova, Michele Tizzani, Daniela Paolotti, and André Panisson. 2022. Echoes through Time: Evolution of the Italian COVID-19 Vaccination Debate. *Proceedings of the International AAAI Conference on Web and Social Media (ICWSM)* (2022).
- [8] Cliona Curley, Eugenia Siapera, and Joe Carthy. 2022. Covid-19 protesters and the far right on Telegram: Co-conspirators or accidental bedfellows? *Social Media + Society* 8, 4 (2022).
- [9] Marco Di Giovanni, Francesco Pierri, Christopher Torres-Lugo, and Marco Brambilla. 2022. VaccinEU: COVID-19 Vaccine Conversations on Twitter in French, German and Italian. *Proceedings of the International AAAI Conference on Web and Social Media (ICWSM)* (2022).
- [10] Fernanda M Dias, Bruno P Rodrigues, Luiz G Araújo, and Ana MS Soares. 2020. Social media influence in the COVID-19 pandemic. *Brazilian Journal of Urology* 46, 6 (2020).
- [11] Gabriele Donzelli, Giacomo Palomba, Ileana Federigi, Francesco Aquino, Lorenzo Cioni, Marco Verani, Annalaura Carducci, and Pierluigi Lopalco. 2018. Misinformation on vaccination: A quantitative analysis of YouTube videos. *Human Vaccines & Immunotherapeutics* 14, 7 (2018).
- [12] John Elias. 2021. YouTube to Add Labels to Some Health Videos amid Misinformation Backlash. <https://www.cnbc.com/2021/07/19/youtube-labeling-some-health-videos-amid-misinformation-backlash.html>.
- [13] Benton Institute for Broadband & Society. 2017. *The Future of Free Speech, Trolls, Anonymity, and Fake News Online*. <https://www.benton.org/headlines/future-free-speech-trolls-anonymity-and-fake-news-online>
- [14] Federico Germani and Nikola Biller-Andorno. 2021. The anti-vaccination infodemic on social media: A behavioral analysis. *PLoS One* 16, 3 (2021).
- [15] Tamar Ginossar, Iain J Cruickshank, Elena Zheleva, Jason Sulskis, and Tanya Berger-Wolf. 2022. Cross-platform spread: Vaccine-related content, sources, and conspiracy theories in YouTube videos shared in early Twitter COVID-19 conversations. *Human Vaccines & Immunotherapeutics* 18, 1 (2022).

- [16] Sharath Chandra Guntuku, Alison M Buttenheim, Garrick Sherman, and Raina M Merchant. 2021. Twitter discourse reveals geographical and temporal variation in concerns about COVID-19 vaccines in the United States. *Vaccine* 39, 30 (2021).
- [17] Eslam Hussein, Prerna Juneja, and Tanushree Mitra. 2020. Measuring misinformation in video search platforms: An audit study on YouTube. *Proceedings of the ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW)* (2020).
- [18] D. Jemielniak and Y. Krempovych. 2021. An analysis of AstraZeneca COVID-19 vaccine misinformation and fear mongering on Twitter. *Public Health* 200 (2021).
- [19] Julie Jiang, Emily Chen, Shen Yan, Kristina Lerman, and Emilio Ferrara. 2020. Political polarization drives online conversations about COVID-19 in the United States. *Human Behavior and Emerging Technologies* 2, 3 (2020).
- [20] Julie Jiang, Xiang Ren, Emilio Ferrara, et al. 2021. Social media polarization and echo chambers in the context of COVID-19: Case study. *JMIRx Med* 2, 3 (2021).
- [21] Garimella Kiran, De Francis Morales Gianmarco, Gionis Aristides, and Mathioudakis Michael. 2018. Quantifying Controversy on Social Media. *ACM Transactions on Social Computing* (2018).
- [22] Aleks Knuutila, Aliksandr Herasimenka, Hubert Au, Jonathan Bright, Rasmus Nielsen, Philip N Howard, et al. 2020. COVID-related misinformation on YouTube: The spread of misinformation videos on social media and the effectiveness of platform policies. *COMPROM Data Memo 2020.6* (2020).
- [23] Nandita Krishnan, Jiayan Gu, Rebekah Tromble, and Lorien C Abrams. 2021. Research note: Examining how various social media platforms have responded to COVID-19 misinformation. *Harvard Kennedy School (HKS) Misinformation Review* 2, 6 (2021), 1–25.
- [24] Valerio La Gatta, Luca Luceri, Francesco Fabbri, and Emilio Ferrara. 2023. The interconnected nature of online harm and moderation: Investigating the cross-platform spread of harmful content between YouTube and Twitter. In *Proceedings of the 34th ACM Conference on Hypertext and Social Media (HT)*.
- [25] Jacopo Lenti, Yelena Mejova, Kyriaki Kalimeri, André Panisson, Daniela Paolotti, Michele Tizzani, and Michele Starnini. 2023. Global misinformation spillovers in the vaccination debate before and during the COVID-19 pandemic: Multilingual Twitter study. *JMIR Infodemiology* 3 (2023).
- [26] Heidi Oi-Yee Li, Adrian Bailey, David Huynh, and James Chan. 2020. YouTube as a source of information on COVID-19: A pandemic of misinformation? *BMJ Global Health* 5, 5 (2020).
- [27] Sahil Loomba, Alexandre de Figueiredo, Simon J Piatek, Kristen de Graaf, and Heidi J Larson. 2021. Measuring the impact of COVID-19 vaccine misinformation on vaccination intent in the UK and USA. *Nature Human Behaviour* 5, 3 (2021).
- [28] Luca Luceri, Stefano Cresci, and Silvia Giordano. 2021. Social media against society: Information manipulation in the 2020 election. In *The Internet and the 2020 election*, Jody Baumgartner and Terri Towner (Eds.), 3–23.
- [29] Ekaterina Malova. 2021. Understanding online conversations about COVID-19 vaccine on Twitter: vaccine hesitancy amid the public health crisis. *Communication Research Reports* 38, 5 (2021), 346–356.
- [30] Shahan Ali Memon and Kathleen M. Carley. 2020. Characterizing COVID-19 misinformation communities using a novel Twitter dataset. *arXiv:2008.00791* (2020).
- [31] Panagiotis Metaxas, Eni Mustafaraj, Kily Wong, Laura Zeng, Megan O'Keefe, and Samantha Finn. 2015. What do retweets indicate? Results from user survey and meta-review of research. *Proceedings of the International AAAI Conference on Web and Social Media (ICWSM)* (2015).
- [32] NewsGuard. 2023. *Rating Process & Criteria*. NewsGuard. <https://www.newsguardtech.com/ratings/rating-process-criteria/>
- [33] Lynnette Hui Xian Ng and Jia Yuan Loke. 2020. Analyzing public opinion and misinformation in a COVID-19 Telegram group chat. *IEEE Internet Computing* 25, 2 (2020), 84–91.
- [34] Gianluca Nogara, Francesco Pierri, Stefano Cresci, Luca Luceri, Petter Törnberg, and Silvia Giordano. 2023. Toxic Bias: Perspective API misreads German as more toxic. *arXiv:2312.12651* (2023).
- [35] Gianluca Nogara, Padinjaredath Suresh Vishnuprasad, Felipe Cardoso, Omran Ayoub, Silvia Giordano, and Luca Luceri. 2022. The Disinformation Dozen: An Exploratory Analysis of Covid-19 Disinformation Proliferation on Twitter. In *The 14th ACM Web Science Conference (WebSci)*. 348–358.
- [36] Francesco Pierri. 2020. The diffusion of mainstream and disinformation news on Twitter: The case of Italy and France. In *Companion Proceedings of The Web Conference (WWW)*. 617–622.
- [37] Francesco Pierri, Matthew R DeVerna, Kai-Cheng Yang, David Axelrod, John Bryden, and Filippo Menczer. 2023. One Year of COVID-19 Vaccine Misinformation on Twitter: Longitudinal Study. *Journal of Medical Internet Research* 25 (2023).
- [38] Reuters. 2019. NewsGuard to help advertisers avoid misinformation. *Reuters* (2019). <https://www.reuters.com/article/us-media-newsguard-idUSKCN1PQ5FV>
- [39] Rafael Santiago, Saira Ijaz, Mohamed El-Sayed, and Amir Mian. 2022. The COVID-19 vaccine social media infodemic: Healthcare providers' missed dose in addressing misinformation and vaccine hesitancy. *Human Vaccines & Immunotherapeutics* 18, 1 (2022).
- [40] Serena Tardelli, Leonardo Nizzoli, Maurizio Tesconi, Mauro Conti, Preslav Nakov, Giovanni Da San Martino, and Stefano Cresci. 2023. Temporal dynamics of coordinated online behavior: Stability, archetypes, and influence. *arXiv:2301.06774* (2023).
- [41] European Union. 2023. *The common EU response to COVID-19*. European Union. https://european-union.europa.eu/priorities-and-actions/common-eu-response-covid-19_en
- [42] Giacomo Villa, Gabriella Pasi, and Marco Viviani. 2021. Echo chamber detection and analysis: A topology- and content-based approach in the COVID-19 scenario. *Social Network Analysis and Mining* 11, 1 (2021), 78.
- [43] Padinjaredath Suresh Vishnuprasad, Gianluca Nogara, Felipe Cardoso, Silvia Giordano, Stefano Cresci, and Luca Luceri. 2024. Tracking fringe and coordinated activity on Twitter leading up to the US Capitol attack. In *The 18th International AAAI Conference on Web and Social Media (ICWSM)*.
- [44] Emily L Wang, Luca Luceri, Francesco Pierri, and Emilio Ferrara. 2023. Identifying and characterizing behavioral classes of radicalization within the QAnon conspiracy on Twitter. In *Proceedings of the International AAAI Conference on Web and Social Media (ICWSM)*.
- [45] Kai-Cheng Yang, Francesco Pierri, Pik-Mai Hui, David Axelrod, Christopher Torres-Lugo, John Bryden, and Filippo Menczer. 2021. The COVID-19 Infodemic: Twitter versus Facebook. *Big Data & Society* 8, 1 (2021).
- [46] Wang Yuxi, McKee Martin, Torbica Aleksandra, and Stuckler David. 2019. Systematic Literature Review on the Spread of Health-related Misinformation on Social Media. *Social Science & Medicine* (2019).