



## City Research Online

### City, University of London Institutional Repository

---

**Citation:** Santoro, A., Galeazzi, A., Scantamburlo, T., Baronchelli, A., Quattrocioni, W. & Zollo, F. (2023). Analyzing the changing landscape of the Covid-19 vaccine debate on Twitter. *Social Network Analysis and Mining*, 13(1), 115. doi: 10.1007/s13278-023-01127-3

This is the published version of the paper.

This version of the publication may differ from the final published version.

---

**Permanent repository link:** <https://openaccess.city.ac.uk/id/eprint/31436/>

**Link to published version:** <https://doi.org/10.1007/s13278-023-01127-3>

**Copyright:** City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

**Reuse:** Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

---

---

---

City Research Online:

<http://openaccess.city.ac.uk/>

[publications@city.ac.uk](mailto:publications@city.ac.uk)

---



# Analyzing the changing landscape of the Covid-19 vaccine debate on Twitter

Arnaldo Santoro<sup>1</sup> · Alessandro Galeazzi<sup>1</sup> · Teresa Scantamburlo<sup>1</sup> · Andrea Baronchelli<sup>2</sup> · Walter Quattrociocchi<sup>3</sup> · Fabiana Zollo<sup>1,4</sup>

Received: 10 March 2023 / Revised: 7 June 2023 / Accepted: 31 August 2023  
© The Author(s) 2023

## Abstract

The issue of vaccine hesitancy has posed a significant challenge during the Covid-19 pandemic, as it increases the risk of undermining public health interventions aimed at mitigating the spread of the virus. While the swift development of vaccines represents a remarkable scientific achievement, it has also contributed to skepticism and apprehension among some populations. Against this backdrop, the suspension of the AstraZeneca vaccine by the European Medicines Agency further exacerbated an already contentious debate around vaccine safety. This paper examines the Twitter discourse surrounding Covid-19 vaccines, focusing on the temporal and geographical dimensions of the discussion. Using over a year's worth of data, we study the public debate in five countries (Germany, France, UK, Italy, and the USA), revealing differences in the interaction structure and in the production volume of questionable and reliable sources. Topic modeling highlights variations in the perspectives of reliable and questionable sources, but some similarities across nations. Also, we quantify the effect of vaccine announcement and suspension, finding that only the former had a significant impact in all countries. Finally, we analyze the evolution of the communities in the interaction network, revealing a relatively stable scenario with a few considerable shifts between communities with different levels of reliability. Our results suggest that major external events can be associated with changes in the online debate in terms of content production and interaction patterns. However, despite the AZ suspension, we do not observe any noticeable changes in the production and consumption of misinformation related to Covid-19 vaccines.

**Keywords** Covid-19 · Vaccine · Misinformation · Infodemic · Twitter · Social media

## 1 Introduction

Public debate is vital to the correct functioning of a democratic society. With the rise of the Internet, digital platforms increased the opportunity for users to engage in public discussions (Flaxman et al. 2016; Schäfer et al. 2015), making this exchange more immediate and

disintermediated (Del Vicario et al. 2016; Quattrociocchi et al. 2014). Online discussions can play a crucial role in the public discourse, influencing policy decisions, and mobilizing action on a variety of issues, such as in the case of political elections (Bovet and Makse 2019; Cinelli et al. 2020). However, the online environment is not a stand-alone entity but is intimately connected with the offline world. Indeed, events such as wars, political elections, epidemic outbreaks, and natural disasters can attract the users' attention and shape online debates (Chen and Ferrara 2022; Bovet and Makse 2019; Cinelli et al. 2020; Sloggy et al. 2021). The outbreak of the coronavirus generated so much hype in online and traditional media that the World Health Organization used the term 'infodemic' to express concern about the excessive amount of information being circulated during the pandemic (Briand et al. 2021).

Vaccines have been the subject of an enduring public discourse that continues to draw widespread attention, as seen

✉ Fabiana Zollo  
fabiana.zollo@unive.it

<sup>1</sup> Ca' Foscari University of Venice, Via Torino, 155,  
30172 Venice, Italy

<sup>2</sup> City University of London, London EC1V 0HB, UK

<sup>3</sup> Sapienza University of Rome, Viale Regina Elena, 295,  
Rome 00161, Italy

<sup>4</sup> The New Institute Centre for Environmental Humanities,  
Dorsoduro 3911, Venice 30123, Italy

recently with the case of Covid-19 (Burki 2019). Vaccine hesitancy and the raise of anti-vax movements may become a public health issue, especially during times of crisis. In some cases, countries strengthened vaccination programs and introduced laws to make certain vaccinations mandatory (Maltezou et al. 2019; MacDonald et al. 2018; Siciliani et al. 2020). The Covid-19 outbreak represents a unique scenario in vaccination history. The unprecedented impact of the pandemic prompted research to accelerate to such an extent that vaccines were developed in an exceptionally short time. The announcements of the first vaccines against Covid-19—first Pfizer and Moderna and then AstraZeneca—generated much debate on online platforms. The European Medical Agency’s (EMA) decision to temporarily suspend the use of a specific batch of the AstraZeneca vaccine as a precautionary measure to evaluate potential issues and side effects has garnered additional attention.

The online debate on vaccines has been intensively studied in correspondence with the Covid-19 outbreak. The authors (Yousefinaghani et al. 2021) analyze 4 million tweets about the vaccine debate over 1 year, studying the prevalence and evolution of the sentiments and opinions. They found that positive sentiment was slightly dominating and attracted higher engagement, but the volume of discussion on vaccine rejection and hesitancy was higher than the one on the interest in vaccines. However, when focusing only on ‘AstraZeneca,’ negative information was the most retweeted content and they often pointed to sources known to be misinformation spreaders, as shown in (Jemielniak and Kremповych 2021). Additionally, the authors emphasize the existence of coordination networks playing a role in political astroturfing and vaccine diplomacy. By utilizing sentiment analysis, the authors of (Aljedaani et al. 2022) analyzed the attitudes and apprehensions toward vaccination among the Arab population on Twitter, revealing variations in their vaccine preferences. The authors of (Mittal et al. 2021) conducted a study to investigate the relationship between online sentiment and disease outbreaks in terms of deaths, infections, and recoveries. By bridging online and offline data, they found a positive association between global rates of infections, deaths, and recoveries and the prevalence of tweets expressing negative sentiment.

In this paper, we analyze how the online public debate evolves in conjunction with significant external events. In particular, here we consider two major events:

1. The announcement of the first Covid-19 vaccine by Pfizer <sup>1</sup>;

2. The suspension of the AstraZeneca (AZ) vaccine by the European Medicines Agency. <sup>2</sup>

We analyzed tweets from five countries: France, Germany, the UK, Italy, and the USA. Each country had its own unique societal context but faced challenges with vaccine hesitancy before the Covid-19 pandemic. For instance, Italy witnessed a negative trend in vaccination coverage, which led to the implementation of compulsory vaccination in July 2017 (Gualano et al. 2018). Prior to the pandemic, France had one of the highest rates of vaccine hesitancy globally (Ward et al. 2019). Similarly, Germany has faced a relatively high level of vaccine skepticism, and surveys indicate a decline in enthusiasm for Covid-19 vaccines throughout the pandemic, despite the population’s strong trust in institutions and health experts (Fiske et al. 2022). In Germany and Italy, concerns regarding the AstraZeneca vaccine, which was ultimately suspended by EMA, contributed to erode trust in scientific and political authorities (Zimmermann et al. 2023). However, the UK, while also facing declining vaccination rates (Kennedy 2020), was not under the jurisdiction of the EMA when it suspended the AstraZeneca vaccine. Moreover, the UK had one of the fastest vaccine roll-outs in Europe (Gallardo 2021). Lastly, the USA, being geographically distant and independent from the other countries considered, was less likely to be affected by the EMA’s decision. However, vaccine hesitancy in the USA has been growing in recent times to the extent that diseases preventable through vaccination, such as measles, have experienced repeated outbreaks due to reduced immunity (Yasmin et al. 2021). Moreover, there has been a decline in childhood vaccination rates, particularly in states where legislation permits personal choice exemptions, indicating a growing vaccine hesitancy in those areas (Lo and Hotez 2017).

First, we compare the interaction structure, content volume, and the most debated topics in five countries (France, Germany, UK, Italy, and the USA) distinguishing between questionable and reliable content. Then, we analyze the temporal evolution of the debate around vaccines and fit auto-regressive models to study how the online public discourse developed in response to significant external events. Our study makes a threefold contribution and addresses the following research questions: (1) What user communities emerge in each country, and how do they compare across countries? (2) What are the characteristics of questionable and reliable narratives in each country, and how do they compare across countries? (3) How has the public debate on vaccination evolved over time in each country?

<sup>1</sup> “Pfizer and BioNTech Conclude Phase 3 Study of Covid-19 Vaccine Candidate, Meeting All Primary Efficacy Endpoints“, published on 18/11/2020 on Pfizer’s website: <https://rb.gy/9wyipm>.

<sup>2</sup> “Covid-19 Vaccine AstraZeneca: PRAC preliminary view suggests no specific issue with batch used in Austria“, published on 10/03/2021 on EMA’s website: <https://rb.gy/0pfbti>.

Our results indicate that the shares of reliable and questionable content can fluctuate based on the country under consideration. Topic modeling reveals different narratives of the debate for content pointing to questionable and reliable sources, although a certain level of similarity across the countries can be found. The temporal analysis shows that the first event (vaccine announcement) influenced the consumption of both questionable and reliable content more than the EMA's suspension of the AZ vaccine batch, although some differences can be found across countries. Finally, we found that users split into two communities rather stable over time based on their attitude toward vaccines, with a small fraction of users changing their beliefs.

This paper contributes to understanding the users' news consumption and reactions to external events in different countries and thus, on the one hand, it provides useful insight into how to communicate negative events during a crisis, on the other hand, it highlights the need for specific communication policy based on the audience characteristics.

## 2 Methods

### 2.1 Data collection

Data was collected from Twitter in the period from January 1st, 2020 to April 30th, 2021. The data collection process was carried out using the Twitter API,<sup>3</sup> which is publicly available, through the full-archive historic search endpoint and academic research product track. According to the API specification, only publicly available data in compliance with users' privacy settings can be accessed. We collected tweets containing an URL and at least one keyword related to Covid-19 vaccines, such as 'immune' and 'pharma' (see Table 1). The final dataset consists in 3, 797, 305 tweets published by 985, 523 accounts.

### 2.2 Tweet classification

For each tweet, we derived three different attributes: the language, the trust score, and the country of origin. The language was automatically detected by Twitter. The trust score was retrieved from NewsGuard, a tool that provides trust ratings for news and information websites on the basis of nine journalistic criteria. These criteria are individually assessed and then combined to produce a single 'trust score' from 0 to 100 for a given news media outlet. The scores are assigned by a team of journalists. Scores are not given to platforms (for example, Twitter and Facebook), individuals, or satire content. The criteria evaluate basic practices of credibility

**Table 1** Vaccine-related keywords for data collection

Country	Vaccine keywords	Brand keywords
France	immun*, vaccin*, dos*, pharma*, médicament*, no-vax*, novax*, pro-vax*, provax*, antivax*, anti-vax*	AstraZeneca, BioNTech, CanSino, CureVac, Jassen,
Germany	immun*, pharma*, vakz*, impf*, no-vax*, novax*, pro-vax*, provax*, antivax*, anti-vax*	Johnson and Johnson, Moderna, Novavax, Pfizer, Sinopharm
Italy	immun*, farma*, pharma*, dos*, vacc*, no-vax*, novax*, pro-vax*, provax*, antivax*, anti-vax*	
UK and US	vaccin*, dose*, pharma*, immun*, no-vax*, novax*, pro-vax*, provax*, antivax*, anti-vax*	

List of vaccine-related terms used for data collection. The asterisk indicates that the term was considered in all its endings

and transparency of news sources.<sup>4</sup> Here, we use NewsGuard's trust scores to distinguish between *questionable* and *reliable* tweets. According to NewsGuard documentation, a source is classified as questionable if its trust score is below 60, while it is considered reliable if the score is equal to or above 60. Therefore, every tweet was categorized according to the trust score of the source (URL) it referenced.

NewsGuard provides trust scores for information sources in five countries, including Germany, France, the UK, Italy, and the USA. Despite the ease of determining the country of origin for a source, assigning a country to a tweet is not always a straightforward process due to the lack of precise geolocalized data. This challenge is particularly pronounced when analyzing English language content, given its global usage. Moreover, we faced the challenge of determining how to assign tweets that mention sources originating from multiple countries.

To address these issues, we implemented a strategy of assigning tweets to a country only if the language and source information were consistent. By doing so, we sought to minimize any ambiguity that may arise from tweets containing mixed-language content or sources from multiple countries. For example, consider a tweet in English language referring to multiple sources in France, UK, and the USA. According to our strategy, the tweet is considered part of the debate in both UK and USA, but not in France. A robustness test was performed using alternative strategies, and no significant

<sup>3</sup> Twitter Developer Platform: <https://developer.twitter.com/>.

<sup>4</sup> More detail regarding the rating process is available at <https://www.NewsGuardtech.com/ratings/rating-process-criteria/>

difference emerged.<sup>5</sup> Additionally, to mitigate potential biases in the keyword selection toward either reliable or questionable news outlets, we computed the percentage of each keyword within both categories (see SI for details). The analysis revealed some minor variations between the two categories, particularly in terms of the frequency of vaccine brand and name mentions. However, these differences do not compromise the validity of our study.

### 2.3 Network analysis

To investigate the presence of segregated communities, we built and analyzed the retweet networks for each country. In particular, we built a weighted undirected network  $N = (V, E, w)$  with vertices  $V$  that represent the accounts retweeting or being retweeted at least once, where two vertices  $u, v \in V$  are connected with an edge  $(u, v) = e \in E$  if and only if  $u$  retweeted  $v$  or  $v$  retweeted  $u$ . We assigned a weight  $w(e)$  to each edge  $e = (u, v)$  corresponding to the number of retweets between  $u$  and  $v$  and a trust score to each vertex  $v \in V$  corresponding to the average trust score of its tweets. Then, We employed the Louvain clustering algorithm (Blondel et al. 2008) to identify groups of accounts that exhibit greater connectivity with each other compared to the rest of the network.

To assess the evolution of communities and the possible changes in the consumption of questionable content, we employed dynamic clustering techniques. For each country, we created a Dynamic Network  $DN = (N_1, N_2, N_3)$ , each DN consisting of three ‘snapshots’  $N_t, t = 1, 2, 3$ . The three snapshots were obtained by considering the retweets produced in the following time intervals: from January 1st 2020 until the Pfizer announcement ( $t = 1$ ), from Pfizer announcement until EMA announcement ( $t = 2$ ), and from the EMA announcement until April 30th 2021 ( $t = 3$ ). On each  $DN$  we employed smoothed Louvain dynamic community detection (Aynaoud and Guillaume 2010), which performs Louvain clustering on each snapshot, using the clustering on the previous slice as initialization. We used the default Jaccard similarity as the matching function and the default matching threshold of 0.3 (Greene et al. 2010).

### 2.4 Topic modeling

To better understand the narratives of these communities, we considered the original tweets and extracted topics using BERTopic (Grootendorst 2022), a state-of-the-art topic modeling tool that extracts latent topics from a collection of documents. We removed URLs from the text and embedded the documents using 768-dimensional pretrained

sentence classification embeddings. Afterward, we reduced their dimensionality to 5 through UMAP, using 15 neighbors and cosine similarity. Finally, we clustered them with HDBSCAN, employing Euclidean distance, and minimum cluster size of  $\max(10, \frac{N}{10^3})$ , where  $N$  is the number of tweets to be clustered. Topic names were assigned by manually inspecting the most representative tweets provided by BerTopic for each cluster of tweets.

### 2.5 Time-series analysis

To study the evolution of content production from reliable and questionable sources over time, we quantified the changes in the production of questionable and reliable tweets in conjunction with the first Covid-19 vaccine announcement by Pfizer (09/11/2020) and the AstraZeneca suspension by EMA (09/03/2021). To do so, we fitted a model to questionable and reliable time series for each country, resulting in a total of ten time series. The model was designed to capture the underlying structural patterns of tweet counts in response to the events. We assumed the time series to be independent of each other and applied a Box–Cox transformation with  $\lambda = 0$  to stabilize the variance of the dependent variable  $Y_t$ . Autocorrelation in the series was addressed by employing dynamic modeling and fitting a linear model with an ARIMA ( $p, q, r$ ) error term. This resulted in the model

$$Y_t^{(\lambda)} = \beta_0 + \beta_1 PFZ_t + \beta_2 EMA_t + \eta_t$$

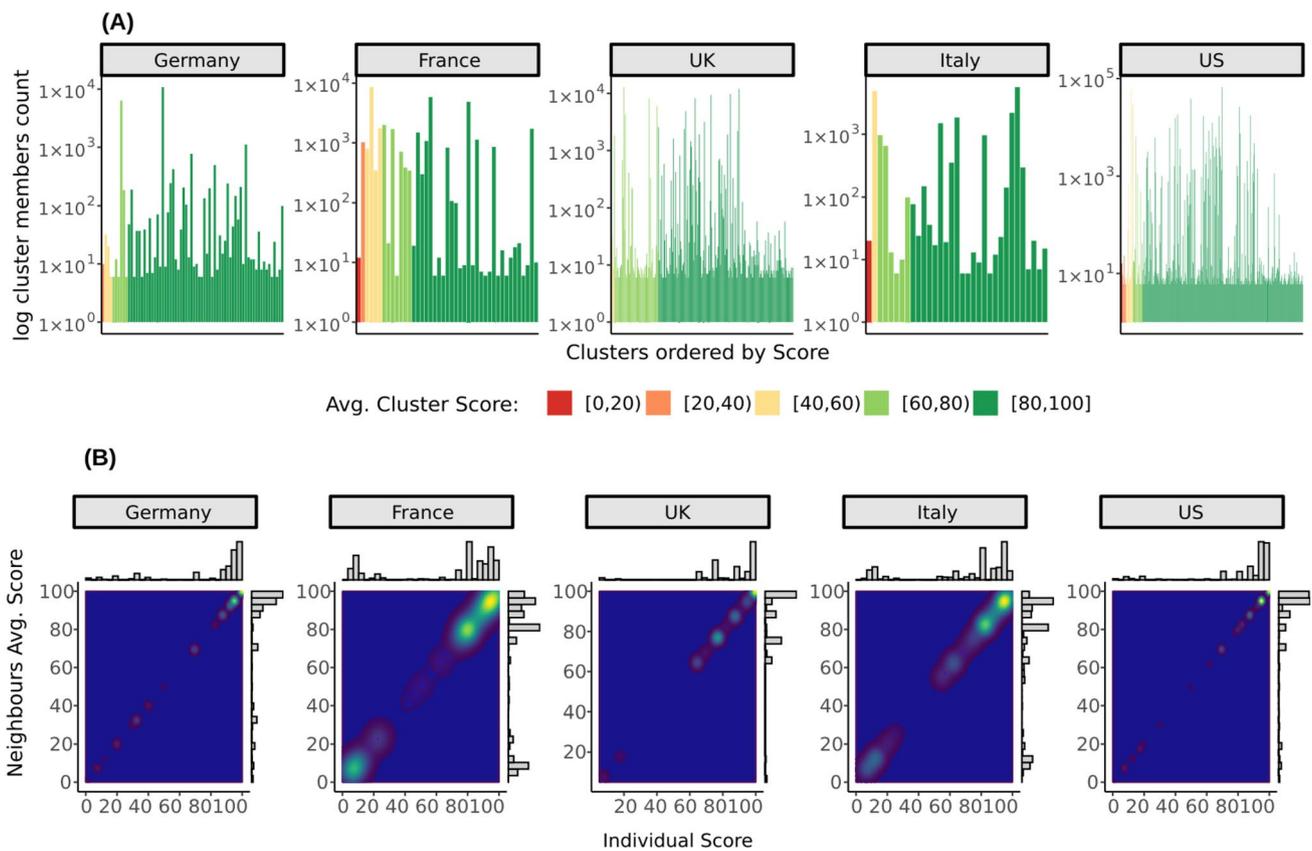
where  $Y_t^{(\lambda)}$  indicates the Box–Cox transformation of the original time series  $Y$  at time  $t$ . Changes in average tweet counts are captured by  $PFZ_t$  and  $EMA_t$ , indicator variables which are set to 0 before the events and 1 otherwise. These changes are deprived of fluctuations in the time series, which in turn are captured by  $\eta_t$ , an error term behaving as an ARIMA( $p, q, r$ ) process that optimizes BIC (Hyndman and Khandakar 2008). The relative change  $\Delta X = \frac{X_{t_2} - X_{t_1}}{X_{t_1}}$  (Table 2)

measures how much the expected tweet production of  $X \in \{Q, R\}$  changed from  $t_1$  to  $t_2$ , where  $Q$  and  $R$  stand for questionable or reliable content, and  $X_{t_1}$  and  $X_{t_2}$  refer to the volume of  $X$  tweets produced until time  $t_1$  and  $t_2$ , respectively.

### 2.6 Limitations

While our analysis offers valuable insights and is based on robust and widely utilized techniques, it is important to acknowledge certain limitations. Firstly, our estimation of user reliability is based on the sources they share. Although this approach is scalable and commonly employed in social media analysis, it may misclassify users who share articles with satirical or critical intents, as it does not consider the broader contextual information.

<sup>5</sup> For details, see Supplementary Information.



**Fig. 1** Clustering results and neighborhood analysis. **A** Result of Louvain clustering on retweet interaction network as a function of trust score. For each cluster, the average score of its users is computed and color-coded. Only clusters with at least five members are shown. **B**

Joint distribution density of individual trust score (x-axis) against neighbors trust score (y-axis) of retweet interaction networks with marginal distributions

Nonetheless, this approach has demonstrated its reliability, enabling consistent analysis, and aligns with similar techniques employed in related studies (Cinelli et al. 2021; Flamino et al. 2023). Secondly, in classifying news domains, we relied on data from NewsGuard, which may have varying coverage in each country. However, NewsGuard rates all sources that account for (at least) 95% of online engagement in each country. Thirdly, in the topic modeling analysis, we assigned names to each cluster of documents manually. To minimize arbitrariness, we began with the most representative documents identified by BerTopic, following the methodology of previous studies (Falkenberg et al. 2022). Lastly, while our set of keywords encompasses a wide range of topics related to Covid-19 vaccines, it is possible that we may have missed some niche keywords. Nevertheless, as mentioned in Sect. 2, our dataset appears to be balanced with respect to different types of sources.

### 3 Results and discussion

#### 3.1 User communities

We first look at the differences and similarities in the structure of the vaccine debate across countries relying on interaction networks (see Sect. 2.3 for details).

In Panel A of Fig. 1, the results of Louvain clustering on retweet networks for each country are presented. Each bar represents a cluster found in the network and its average trust score is color-coded. The clusters are arranged in ascending order of average trust score. Across all countries, the majority of communities are made on average by highly reliable users, with at least 74% of them having a high average score. However, a small number of clusters with low average scores are present in all countries.



**Fig. 2** Most debated topics by country. The five most frequent topics, identified through topic modeling, by country and reliability. Red (green) panels show topics in questionable (reliable) tweets. Lengths

of the bars represent the share of tweets belonging to each topic relative to the other four topics. The percentage indicates the share of tweets by country and reliability belonging to each topic

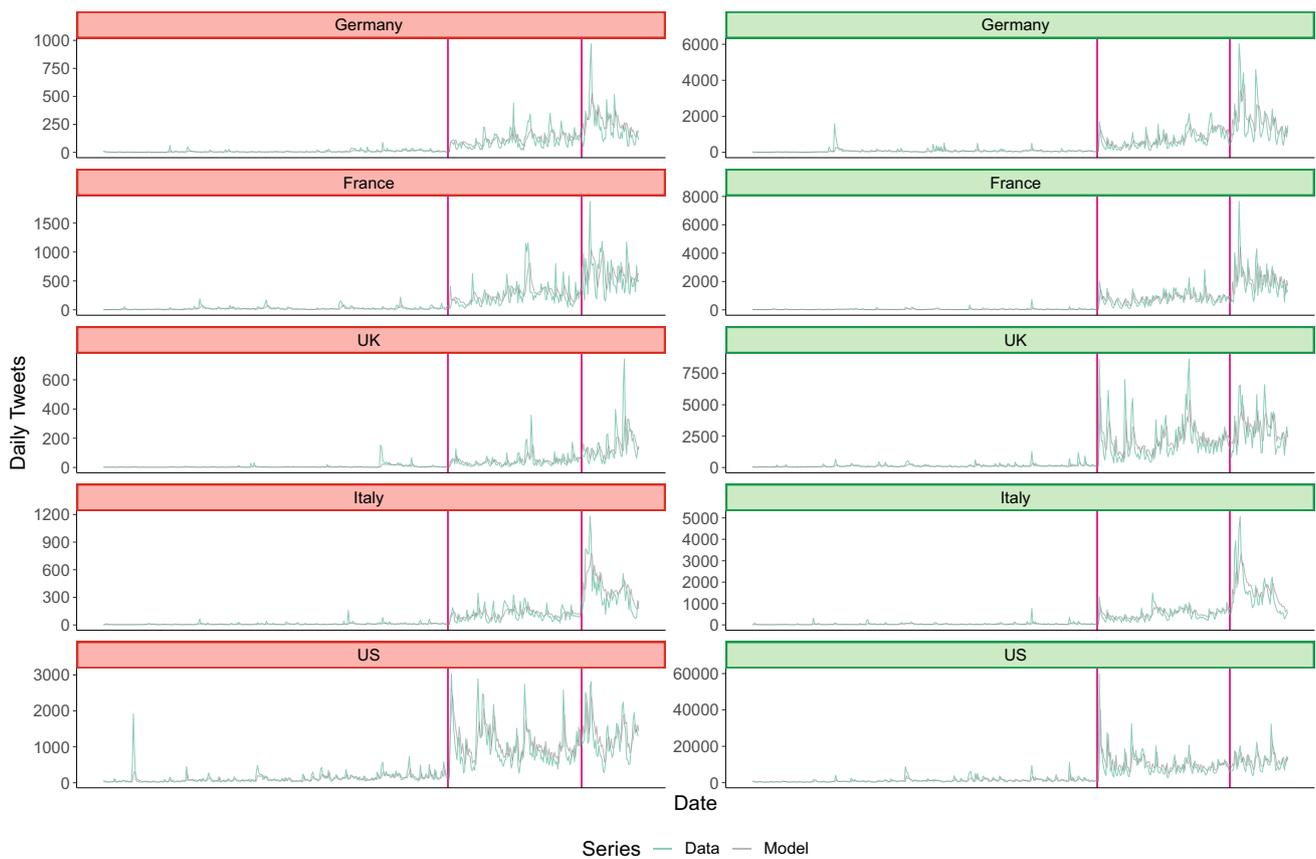
Panel B of Fig. 1 reports the joint distribution of the accounts' individual trust score against the average trust score of their neighbors. We observe the presence of separated communities laying on the diagonal, indicating that users tend to interact with others sharing sources with similar trust scores. The observed pattern supports the presence of an echo chamber effect, for which users tend to interact with others who hold similar views and rarely engage with those who have opposing perspectives.

### 3.2 Questionable and reliable narratives

To better understand the topics of discussion in each country for questionable and reliable sources, we conducted topic modeling analysis (see Sect. 2.3 for details).

Here, we concentrate on the distinctions in the overall narratives conveyed by the two groups, and therefore, we confine our analysis to the static version, without considering the temporal evolution of the debate. In the subsequent subsection, we delve into the temporal dynamics of

the discourse. The results presented in Fig. 2 reveal some differences in the types of topics discussed: while both categories of tweets addressed vaccine side effects, this topic was more frequently discussed among questionable sources compared to reliable sources. Moreover, both types of tweets in each country also touched on political topics related to the government's agreements to obtain, approve, and suspend vaccines. However, the narrative of the discussions differed between the two groups. Questionable sources often expressed concerns about excessive government control, while reliable sources focused on inefficiencies in the vaccination campaign and the lack of clarity in AstraZeneca vaccine recommendations. The topic of vaccine efficacy was also discussed in both groups, but the nature of the discussions differed significantly. While reliable sources mainly addressed the efficacy statistics of the vaccines, questionable sources talked about alternative treatments, signaled by the presence of terms like 'hydroxychloroquine' or 'vitamins.' Furthermore, they questioned the usefulness of vaccines and challenged medical studies on their efficacy. One



**Fig. 3** Volume of content related to Covid-19 Vaccine on Twitter. Red (green) panels show questionable (reliable) daily Tweet counts of each country; the green line shows actual data, and the model's

fitted values are in gray. The first vertical red lines mark the Pfizer announcement, and the second red lines mark the start of the EMA suspension

of the prominent topics discussed by reliable sources was Covid-19 official reports. On the other hand, in questionable sources topics such as fear of side effects (including concerns over ‘facial paralysis,’ ‘DNA damages,’ ‘abortion,’ and ‘fertility risks’) and conspiracy theories (such as ‘Bill Gates’ involvement in the pandemic, ‘illegal experiments on humans,’ ‘sterilization programs,’ and ‘chip implantation’) were more present.

**3.3 Evolution of the public debate over time**

Figure 3 shows the amount of content (number of tweets) produced by questionable and reliable sources over time. We notice that content production increased in correspondence with both the announcement of the Pfizer vaccine (November 9th 2020) and the EMA suspension (March 9th 2021). To better understand the changes in tweet production following the considered events, we fitted a dynamic model to data, represented by the gray line in Fig. 3 (see Sect. 2.3 for details). Our models indicate that there was a statistically significant increase in tweet production of both questionable and reliable sources following the Pfizer

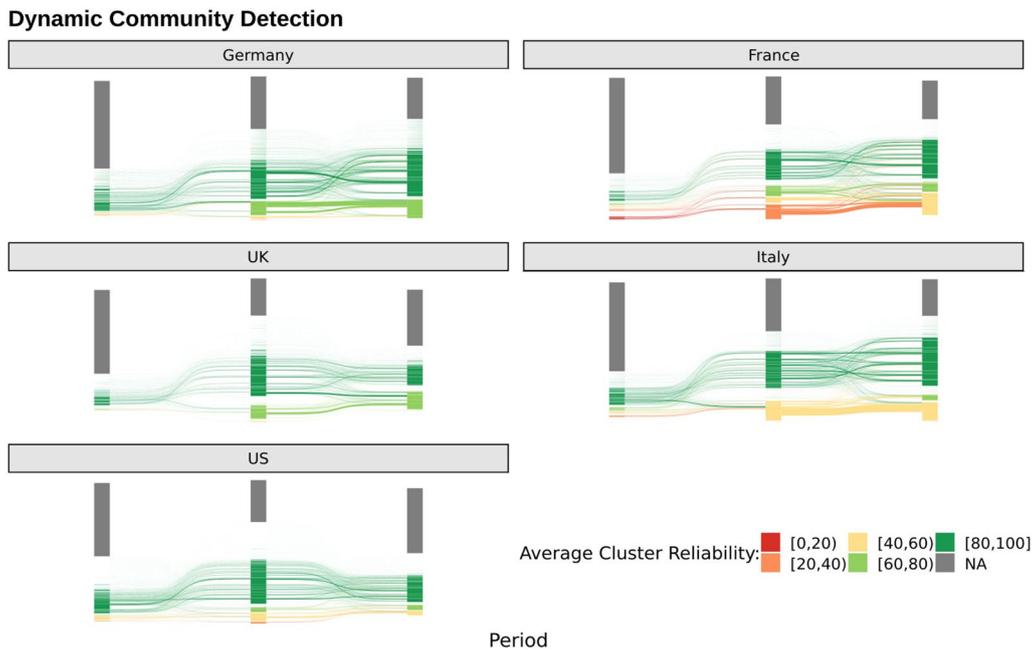
announcement, however, some differences among countries arise when comparing the growth of questionable and reliable content. Indeed, while Germany experienced a higher growth of questionable sources in comparison with reliable ones, in all the other countries, reliable sources grew more than questionable ones. To measure these differences, we calculated the relative changes in content production and compared them by computing the ratio between the changes in questionable content production and those in reliable content production. If the growth rate of questionable sources surpasses that of reliable ones, the resulting ratio will be greater than 1, whereas it will be lower if the opposite is true.

The results presented in Table 2 show that all countries but Germany experienced a greater growth in the number of reliable tweets following the Pfizer vaccine announcement. This can be explained by the different volumes of questionable sources before the event: while Germany had a very small volume of tweets referring to questionable sources (10% of reliable sources volume), in France this presence is higher than in any other country (48% of reliable sources volume).

**Table 2** Relative changes in expected tweet production

Relative changes	Rating	Germany	France	UK	Italy	US
Changes After PFZ	$\Delta Q$	21.38	12.55	18.50	15.55	11.08
	$\Delta R$	15.15	24.30	21.27	19.54	12.51
	$\Delta Q/\Delta R$	1.41	0.52	0.87	0.80	0.89
Changes After EMA	$\Delta Q$	0.84*	1.02*	1.07*	2.29	0.36*
	$\Delta R$	1.36	1.52	0.38*	1.49	0.24*
	$\Delta Q/\Delta R$	0.62	0.67	2.81	1.54	1.47

$\Delta Q$  ( $\Delta R$ ) measures the shift in the volume of the average questionable (reliable) tweet production after an event compared to the level before the event, as measured by the model. The \* indicates statistics employing non-statistically significant ( $p$  value > 0.05) coefficients; i.e., there is not enough evidence in favor of a change in level after the event. The ratio  $\Delta Q/\Delta R$  measures the changes of questionable average tweet production with respect to the reliable one: numbers over 1 indicate that questionable tweets grew more than reliable tweets; conversely, numbers below 1 indicate that reliable grew more



**Fig. 4** Evolution of the communities in the interaction network over time. Communities are represented as lines and connections between communities represent accounts movement. Average community trust

score is color-coded and gray bars represent accounts that do not participate in the period, but do participate in the country’s debate in other periods

This suggests that the questionable debate in France may have experienced a saturation effect due to the already high preexisting volume of questionable content, resulting in a smaller growth in questionable tweets after the event. Conversely, statistics on Germany reveal a lower level of tweet activity that saw the largest growth in the number of questionable tweets. Nonetheless, the volume of questionable sources remained significantly lower than that of reliable sources in all countries.

Next, we examine the effect of the suspension of the AstraZeneca vaccine by EMA, which represents the second event under consideration. The models reveal an increase in both questionable and reliable content, but the magnitude

of this change is significantly lower than that observed during the vaccine announcements. Indeed, the second half of Table 2 shows that the increase is one order of magnitude smaller. Moreover, Italy is the only country that shows statistically significant increases in both questionable and reliable sources. In contrast, neither the UK nor the USA show statistically significant changes in either type of source.

Given that the temporal analysis has revealed significant changes in the prevalence of questionable and reliable sources, a natural question is whether these events are also associated with changes in the community structure and opinion of users. To study the evolution of the communities, we performed Dynamic Community Detection on the

networks. This approach enabled us to track the changes in trust score over time, at both the community and account levels. Figure 4 shows the evolution of the communities during three distinct time spans: before the Pfizer announcement, between the Pfizer announcement and the AstraZeneca suspension, and after the AstraZeneca suspension. The distribution of clusters varies significantly across countries, with France having a sizable community of questionable accounts, while other countries like the UK and Germany have much fewer such accounts, consistent with the findings in the previous sections. Moreover, we can notice that, in the final snapshot, the merging of the most questionable community with a reliable one resulted in a community with an average score in the reliable range, whereas in other countries such as Italy and the UK, the communities appear to be more stable over time. Also, we observe that considering the trust score threshold of 60, the communities tended to remain relatively stable over time, with only a few accounts shifting between them. Overall, we may notice that the communities with a trust score threshold lower than the country's average tended to exhibit greater cohesiveness and stability, while the more reliable communities appeared to be more dynamic and scattered. One possible explanation for this phenomenon could be that accounts within low-trust-score communities tend to be more active and engaged in the debate, while accounts within higher-trust-score communities are characterized by higher turnover rates. This is consistent with findings from other studies, which suggest that polarized users tend to be more active and engaged in debates (Schmidt et al. 2018). Differences across countries in terms of both production and consumption of questionable content, and response to external events might be influenced by a variety of factors, including level of trust in institutions and social environments (Zimmermann et al. 2023; Sturgis et al. 2021).

## 4 Conclusions

In this study, we used Twitter data to examine the debate surrounding Covid-19 vaccines. Our analysis proceeded in two main directions, investigating cross-country variations and tracing the evolution of the conversation in response to major events, such as the announcement of the Pfizer vaccine and the EMA's suspension of a batch of the AstraZeneca vaccine. Moreover, we incorporated third-party data to explore the presence of misinformation over time and employed topic modeling to uncover emergent narratives. Finally, we aimed to shed light on the structure and dynamics of the vaccine debate on Twitter.

Our findings show that while reliable sources dominate the discussion, communities of users who consume

questionable content are present to varying degrees across different countries. Moreover, the narratives from reliable and questionable sources diverge, although some similarities exist across countries. The analysis of the response to the Pfizer vaccine announcement and the AstraZeneca suspension revealed that both reliable and questionable content increased following these events, though the volume increase after the latter was typically lower and not statistically significant. Finally, our community detection analysis revealed a relatively stable scenario, with only a few shifts between communities.

Taken together, our results confirm the polarized nature of the Covid-19 vaccine debate on Twitter, with stable communities and differences across countries. Furthermore, we show that significant external events can be associated with changes in the production and consumption of online information, though we found no evidence of increased misinformation after the EMA's suspension of the AstraZeneca vaccine. These findings may be relevant for policy and decision makers involved in risk and crisis communication. To maximize the effectiveness of public health policy strategies, it is essential to consider the social and cultural environment in which they are implemented (Cascini et al. 2022). To this aim, the dynamic monitoring of public response through online data may prove extremely helpful in adjusting public health measures. This study shows how social media and network analysis can be used to reveal differences and similarities across countries, thereby highlighting its potential to facilitate the development of tailored plans of action and interventions suited to the population's unique context. Further, enhancing our understanding of the narratives prevalent within distinct communities and countries can assist public organizations in delivering more impactful communication, particularly during times of crisis. Expanding our research to include other social media platforms and socially significant topics such as climate change could provide further insights.

**Supplementary Information** The online version contains supplementary material available at <https://doi.org/10.1007/s13278-023-01127-3>.

**Acknowledgements** All authors acknowledge financial support from the IRIS Academic Research Group. FZ acknowledges financial support from the EU REC project EUMEPLAT GA 101004488 and from the project SERICS (PE00000014) under the NRRP MUR program funded by the EU-NGEU. The authors wish to thank Flavio Vella for fruitful discussions.

**Author contributions** AG, AB, WQ, and FZ conceived the research; AS and AG downloaded the data; AS and AG performed the analysis; All authors analyzed the results and wrote the manuscript.

## Declarations

**Conflict of interest** The authors declare no competing interests.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, 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 licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence 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. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Aljedaani W, Abuhaimed I, Rustam F, Mkaouer MW, Ouni A, Jenhani I (2022) Automatically detecting and understanding the perception of Covid-19 vaccination: a middle east case study. *Soc Netw Anal Min* 12(1):128
- Aynaud T, Guillaume JL (2010) Static community detection algorithms for evolving networks. In: 8th international symposium on modeling and optimization in mobile, Ad Hoc, and wireless networks, pp 513–519
- Blondel VD, Guillaume JL, Lambiotte R, Lefebvre E (2008) Fast unfolding of communities in large networks. *J Stat Mech Theory Exp* 2008(10):P10008
- Bovet A, Makse HA (2019) Influence of fake news in Twitter during the 2016 US presidential election. *Nat Commun* 10(1):7
- Briand SC, Cinelli M, Nguyen T, Lewis R, Prybylski D, Valensise CM, Colizza V, Tozzi AE, Perra N, Baronchelli A et al (2021) Infodemics: a new challenge for public health. *Cell* 184(25):6010–6014
- Burki T (2019) Vaccine misinformation and social media. *Lancet Digital Health* 1(6):e258–e259
- Cascini F, Pantovic A, Al-Ajlouni YA, Failla G, Puleo V, Melnyk A, Lontano A, Ricciardi W (2022) Social media and attitudes towards a covid-19 vaccination: a systematic review of the literature. *EClinicalMedicine* 101454
- Chen E, Ferrara E (2022) Tweets in time of conflict: a public dataset tracking the Twitter discourse on the war between Ukraine and Russia. arXiv preprint [arXiv:2203.07488](https://arxiv.org/abs/2203.07488)
- Cinelli M, Cresci S, Galeazzi A, Quattrociochi W, Tesconi M (2020) The limited reach of fake news on Twitter during 2019 European elections. *PLoS One* 15(6):e0234689
- Cinelli M, Morales GDF, Galeazzi A, Quattrociochi W, Starnini M (2021) The echo chamber effect on social media. In: Proceedings of the national academy of sciences 118(9)
- Cinelli M, Quattrociochi W, Galeazzi A, Valensise CM, Brugnoli E, Schmidt AL, Zola P, Zollo F, Scala A (2020) The Covid-19 social media infodemic. *Sci Rep* 10(1):16598. <https://doi.org/10.1038/s41598-020-73510-5>
- Del Vicario M, Bessi A, Zollo F, Petroni F, Scala A, Caldarelli G, Stanley HE, Quattrociochi W (2016) The spreading of misinformation online. *Proc Natl Acad Sci* 113(3):554–559
- Falkenberg M, Galeazzi A, Torricelli M, Di Marco N, Larosa F, Sas M, Mekacher A, Pearce W, Zollo F, Quattrociochi W et al (2022) Growing polarization around climate change on social media. *Nat Clim Change* 12(12):1114–1121
- Fiske A, Schönweitz F, Eichinger J, Zimmermann B, Hangel N, Sierawska A, McLennan S, Buyx A (2022) The Covid-19 vaccine: trust, doubt, and hope for a future beyond the pandemic in Germany. *PLoS One* 17(4):e0266659
- Flamino J, Galeazzi A, Feldman S, Macy MW, Cross B, Zhou Z, Serafino M, Bovet A, Makse HA, Szymanski BK (2023) Political polarization of news media and influencers on twitter in the 2016 and 2020 US presidential elections. *Nat Hum Behav* 1–13
- Flaxman S, Goel S, Rao JM (2016) Filter bubbles, echo chambers, and online news consumption. *Public Opin Q* 80(S1):298–320
- Gallardo C (2021) 8 reasons the uk leads europe's coronavirus vaccination race. *POLITICO Europe*
- Greene D, Doyle D, Cunningham P (2010) Tracking the evolution of communities in dynamic social networks. In: 2010 international conference on advances in social networks analysis and mining. pp 176–183
- Grootendorst M (2022) Bertopic: neural topic modeling with a class-based tf-idf procedure. arXiv preprint [arXiv:2203.05794](https://arxiv.org/abs/2203.05794)
- Gualano M, Bert F, Voglino G, Buttinelli E, D'Errico M, De Waure C, Di Giovanni P, Fantini M, Giuliani A, Marranzano M, Masanotti G, Massimi A, Nante N, Pennino F, Squeri R, Stefanati A, Signorelli C, Siliquini R, Castaldi S, Di Donna F, Di Martino G, Genovese C, Golferà M, Gori D, Greco P, Loperto I, Miduri A, Olivero E, Prospero E, Quattrocchio F, Rossello P, Rosso A, Sisti L, Stracci F, Zappalà G (2018) Attitudes towards compulsory vaccination in italy: results from the Navidad multicentre study. *Vaccine* 36(23):3368–3374. <https://doi.org/10.1016/j.vaccine.2018.04.029>
- Hyndman RJ, Khandakar Y (2008) Automatic time series forecasting: the forecast package for r. *J Stat Softw* 27(3):1–22. <https://doi.org/10.18637/jss.v027.i03>
- Jemielniak D, Krempovych Y (2021) An analysis of Astrazeneca Covid-19 vaccine misinformation and fear mongering on Twitter. *Public Health* 200:4–6. <https://doi.org/10.1016/j.puhe.2021.08.019>
- Kennedy J (2020) Vaccine hesitancy: a growing concern. *Pediatr Drugs* 22(2):105–111. <https://doi.org/10.1007/s40272-020-00385-4>
- Lo NC, Hotez PJ (2017) Public health and economic consequences of vaccine hesitancy for measles in the united states. *JAMA Pediatr* 171(9):887–892
- MacDonald NE, Harmon S, Dube E, Steenbeek A, Crowcroft N, Opel DJ, Faour D, Leask J, Butler R (2018) Mandatory infant & childhood immunization: rationales, issues and knowledge gaps. *Vaccine* 36(39):5811–5818
- Maltezou HC, Botelho-Nevers E, Brantsæter AB, Carlsson RM, Heining U, Hübschen JM, Josefsdottir KS, Kassianos G, Kyncl J, Ledda C et al (2019) Vaccination of healthcare personnel in europe: update to current policies. *Vaccine* 37(52):7576–7584
- Mittal R, Mittal A, Aggarwal I (2021) Identification of affective valence of Twitter generated sentiments during the Covid-19 outbreak. *Soc Netw Anal Min* 11(1):108
- Quattrociochi W, Caldarelli G, Scala A (2014) Opinion dynamics on interacting networks: media competition and social influence. *Sci Rep* 4(1):1–7
- Schäfer MS et al (2015) Digital public sphere. *Int Encycl Polit Commun* 15:1–7
- Schmidt AL, Zollo F, Scala A, Betsch C, Quattrociochi W (2018) Polarization of the vaccination debate on Facebook. *Vaccine* 36(25):3606–3612. <https://doi.org/10.1016/j.vaccine.2018.05.040>
- Siciliani L, Wild C, McKee M, Kringos D, Barry MM, Barros PP, De Maeseneer J, Murauskiene L, Ricciardi W et al (2020) Strengthening vaccination programmes and health systems in the European union: a framework for action. *Health Policy* 124(5):511–518
- Sloggy MR, Suter JF, Rad MR, Manning DT, Goemans C (2021) Changing climate, changing minds? The effects of natural disasters on public perceptions of climate change. *Clim Change* 168:1–26

- Sturgis P, Brunton-Smith I, Jackson J (2021) Trust in science, social consensus and vaccine confidence. *Nat Hum Behav* 5(11):1528–1534
- Ward JK, Peretti-Watel P, Bocquier A, Seror V, Verger P (2019) Vaccine hesitancy and coercion: all eyes on France. *Nat Immunol* 20(10):1257–1259. <https://doi.org/10.1038/s41590-019-0488-9>
- Yasmin F, Najeeb H, Moeed A, Naeem U, Asghar MS, Chughtai NU, Yousaf Z, Seboka BT, Ullah I, Lin CY et al (2021) Covid-19 vaccine hesitancy in the united states: a systematic review. *Front Public Health* 9:770985
- Yousefinaghani S, Dara R, Mubareka S, Papadopoulos A, Sharif S (2021) An analysis of Covid-19 vaccine sentiments and opinions on Twitter. *Int J Infectious Diseases* 108:256–262. <https://doi.org/10.1016/j.ijid.2021.05.059>
- Zimmermann BM, Paul KT, Araújo ER, Buyx A, Ferstl S, Fiske A, Kraus D, Marelli L, McLennan S, Porta V et al (2023) The social and socio-political embeddedness of Covid-19 vaccination decision-making: a five-country qualitative interview study from Europe. *Vaccine* 41(12):2084–2092

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.