

City Research Online

City, University of London Institutional Repository

Citation: Loru, E., Galeazzi, A., Bonetti, A., Sangiorgio, E., Di Marco, N., Cinelli, M., Falkenberg, M., Baronchelli, A. & Quattrociocchi, W. (2025). Ideology and polarization set the agenda on social media. Scientific Reports, 15, 35816. doi: 10.1038/s41598-025-19776-z

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/36211/

Link to published version: https://doi.org/10.1038/s41598-025-19776-z

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

scientific reports



OPEN

Ideology and polarization set the agenda on social media

Edoardo Loru^{1⊠}, Alessandro Galeazzi², Anita Bonetti³, Emanuele Sangiorgio⁴, Niccolò Di Marco⁵, Matteo Cinelli⁶, Max Falkenberg⁷, Andrea Baronchelli⁶, & Walter Quattrociocchi⁶[™]

The abundance of information on social media has reshaped public discussions, shifting attention to the mechanisms that drive online discourse. This study analyzes large-scale Twitter (now X) data from three global debates—Climate Change, COVID-19, and the Russo-Ukrainian War—to investigate the structural dynamics of engagement. Our findings reveal that discussions are not primarily shaped by specific categories of actors, such as media or activists, but by shared ideological alignment. Users consistently form polarized communities, where their ideological stance in one debate predicts their positions in others. This polarization transcends individual topics, reflecting a broader pattern of ideological divides. Furthermore, the influence of individual actors within these communities appears secondary to the reinforcing effects of selective exposure and shared narratives. Overall, our results underscore that ideological alignment, rather than actor prominence, plays a central role in structuring online discourse and shaping the spread of information in polarized environments.

Keywords Social Media, Ideological Alignment, Polarization, Online Discourse, Online Engagement, Public Debates

The rapid expansion of social media has fundamentally changed the way people communicate, share information, and consume content. These platforms are now central to society, serving as tools for dissemination, entertainment, and consumption of information^{1–4}. However, their engagement-driven business models and the dynamics of online interactions have raised concerns about their broader social impacts^{5,6}, including their role in deepening polarization^{7–9}, propagating misinformation^{10–12}, and amplifying hate speech^{13–17}. Moreover, recent research has highlighted the central role that the structure of social networks has in disseminating information¹⁸.

Building on these concerns, today's digital environment is characterized by an overabundance of information, offering users an unprecedented range of choices for media consumption¹⁹. This shift results from the proliferation of various sources of content, fueling intense competition for user attention^{20–22}, amplified by the engagement-driven strategies of social media platforms. Consequently, the information ecosystem has undergone a significant transformation, leading to a fragmented audience willing to participate^{23,24} but dispersed across numerous platforms and personalized feeds²⁵. Traditional media gatekeepers, once central to public discourse, no longer unilaterally set the agenda²⁶, diverging from previously dominant patterns of centralized media consumption^{27–30}.

This radical transformation has led to a vast and unstructured flow of content³¹ that may induce information overload³². Users often feel overwhelmed by the sheer volume of content when seeking information³³, complicating their ability to navigate this landscape. In this information-saturated ecosystem, online users increasingly exhibit selective consumption behaviors, engaging primarily with content aligned with their existing beliefs^{12,34,35} while ignoring opposing viewpoints^{36,37}. This behavior reinforces echo chambers and exacerbates polarization^{8,38,39}, in an environment where the value of news content is no longer determined solely by its intrinsic quality⁴⁰. News consumption has become a shared social experience, with individuals exchanging links and recommendations within their networks and treating news as a form of cultural currency⁴¹. As a result,

¹Department of Computer, Control and Management Engineering, Sapienza University of Rome, Rome, Italy. ²Department of Mathematics, University of Padova, Padova, Italy. ³Department of Communication and Social Research, Sapienza University of Rome, Rome, Italy. ⁴Department of Social Sciences and Economics, Sapienza University of Rome, Rome, Italy. ⁵Department of Legal, Social, and Educational Sciences, Tuscia University, Viterbo, Italy. ⁶Department of Computer Science, Sapienza University of Rome, Rome, Italy. ⁷Department of Network & Data Science, Central European University, Vienna, Austria. ⁸Department of Mathematics, City St George's, University of London, London, UK. ⁹The Alan Turing Institute, London, UK. [∞]email: edoardo.loru@uniroma1.it; walter.quattrociocchi@uniroma1.it

traditional hierarchical models of media influence struggle to explain the dynamics of public discourse in this decentralized environment ^{42–44}, where narratives and participants compete for prominence.

Previous studies have explored how social media have changed the traditional gatekeeping function of legacy media. Contemporary debates on social platforms are frequently shaped not by institutional actors, but by networked publics⁴⁵ participating in the online public sphere. In classic media research, gatekeeping was understood as the process through which mass media filtered, shaped, and prioritized information flows to audiences^{46,47}. However, the rise of participatory digital technologies has profoundly challenged this top-down model. Scholars have argued that social media has empowered users to actively contribute to the creation and dissemination of news content, thereby eroding the media's monopoly on gatekeeping^{48,49}. As Bro and Wallberg⁵⁰ note, individuals increasingly access news through peer networks, such as family and friends, rather than through editorially curated channels. This shift has altered not only how news spreads, but also who holds the power to define it. While news organizations continue to seek authority over informational flows⁵¹, alternative actors have emerged as influential newsmakers in the digital sphere. This transformation has led some scholars to reconceptualize gatekeeping as a distributed, collaborative process within communicative networks. Hoskins and O'Loughlin⁵² initially questioned whether gatekeeping could be meaningfully applied to such networked contexts, but their research—along with subsequent studies—supports the notion of "network gatekeeping," a dynamic process emerging from diffused and collaborative gatekeeping processes^{52,53}. Our findings contribute to this literature by illustrating how agenda-setting processes might be increasingly negotiated among a wide array of actors, in an environment where news also emerges from bottom-up dynamics.

While classical agenda-setting theory, originating with McCombs and Shaw⁵⁴, was built around a top-down logic typical of mass communication, subsequent scholarship has explored its evolution in light of media digitalization, debating the relevance of the theory in digital settings. The same authors, on the 50th anniversary of the theory, traced its evolution by identifying seven distinct facets, emphasizing need for orientation, network agenda setting, and agendamelding as particularly active areas of contemporary research^{55,56} within an increasingly complex and interactive media ecosystem⁵⁷. Some scholars have argued that agenda setting effects remain operative, albeit adapted to the new media landscape⁵⁸, while others have pointed to a weakening of traditional media's influence in favor of digital platforms^{59,60}, or have conceptualized social media themselves as a further development of agenda setting mechanisms⁶¹. Russell Neuman et al.⁶² identified several patterns of issue framing in social and traditional media that highlighted an interdependence of the two rather than a one-way agenda-setting dynamic. Feezell⁶³, instead, investigated how social media contributes to the dissemination of mainstream news and their perceived importance. New studies on agenda setting processes have also kept in consideration various platforms, such as Twitter^{64–66}, Instagram⁶⁷ and Reddit⁶⁸. Analyzing 71.77 million tweets, Yi and Wang⁶⁹ found that professional media influence only a small number of individuals, while they argue that opinion leaders have become strong competitors of traditional professional media in agenda setting processes.

Understanding these processes requires a change in perspective. Using multiple large-scale Twitter (now X) datasets, our analysis outlines the drivers of public discourse by examining three global debates—climate change (COP26), COVID-19, and the Russo-Ukrainian War. We compare various actor types, such as media outlets, activists, and individual influencers, to determine whether specific categories dominate discussions and what drives online debates. Furthermore, our diverse selection of debates, spanning environmental, health, and geopolitical domains, enables us to assess whether the dynamics we investigate are tied to specific topics or reflect systemic patterns across issues.

In particular, we aim to address the following questions: What are the drivers of user engagement in fragmented social media environments? Do actor categories or ideological alignment better explain how attention is allocated? Our findings reveal a fragmented environment where no single category of actor consistently dominates discussions. Furthermore, users cluster into two polarized ideological communities, one representing the majority of users endorsing a mainstream narrative and a minority usually embracing alternative views. Notably, a user's stance in one debate is a reliable marker of their opinion on other topics. For instance, users supporting climate action in the COP26 debate are likely to advocate for vaccination during COVID-19 and aid for Ukraine, while opponents of COVID-19 vaccination are typically climate skeptics and critics of military aid. This consistency in ideological stance across heterogeneous debates suggests that polarization on social media emerges from overarching ideological divides, rather than being confined to specific issues.

Our results underscore that online discussions are shaped by shared ideological alignments, rather than by a top-down model dominated by traditional gatekeepers. This highlights how decentralized, bottom-up dynamics dominate public discourse, where ideological structures are pivotal in predicting the evolution of debates and the spread of specific narratives. By identifying consistent patterns in user engagement and ideological alignment, this study provides insights into the mechanisms shaping public discourse in an information-saturated landscape.

Results and discussion

We investigate the relationship between engagement and ideological configurations across three major debates—COP26, COVID-19, and the Russo-Ukrainian War—employing data collected with specific keyword searches (see Methods). We begin by analyzing the influence of various types of actors in driving online discussions. Next, we quantify the degree of polarization surrounding different topics and examine the relationships between users' opinions across these topics. Finally, we explore how users in different ideological communities interact with various types of actors across topics.

Decentralized engagement: beyond influencer categories

To investigate the dynamics of engagement in online discussions, we identify a set of accounts that we refer to as 'influencers'. We first select the top 1% of accounts receiving the highest number of retweets within each debate. We further refine the selection from this subset to include only the top 50% producers of original tweets. This

approach ensures that our analysis focuses on accounts that receive significant attention and contribute to the discussion by actively creating new content.

Each influencer is then manually assigned to one of six categories: Activist, International Organization/NGO, Media, Politics, Private Individual, and Other. Accounts whose tweets are unrelated to the topics of discussion are excluded from the analysis. In addition to categorizing these accounts according to their role, we annotate each account with their stance on key issues within the debates. The final sets of influencers are comprised of 561 accounts for COP26, 1577 for COVID-19, and 1166 for Ukraine. The labeling process is described more in detail in Methods and in Supplementary Information, where we also report the proportion of influencers in each category (Supplementary Fig. S1) and the distributions of original tweets posted by these accounts (Supplementary Fig. S2).

One fundamental question is the following: according to agenda-setting theory, we might expect media or political accounts to receive higher levels of engagement than other types of information providers. To understand if this is the case, namely whether the engagement that influencers receive depends on the category they belong to, we analyze the distribution of retweets received by each category per debate. This analysis enables us to assess whether certain categories of influencers are more effective in driving users' engagement.

The distributions displayed in Fig. 1 do not indicate that specific categories of actors consistently receive more retweets. Instead, all categories receive comparable attention, suggesting that the category itself may not be a determining factor for the volume of retweets. For instance, although @GretaThunberg stands out as a prominent figure in the COP26 discussion, other activists are retweeted with comparable frequency to influencers from other categories. Similarly, in the Ukraine debate, accounts in the Media category attract engagement levels similar to influencers from the other categories, with few exceptions such as @KyivIndependent—a team of independent journalists.

To statistically test these insights, we perform pairwise comparisons between the categories within each debate using Kolmogorov-Smirnov tests, which assess whether the retweet distributions for any two categories are drawn from the same population. The results, reported in Supplementary Table S1, show no statistically significant differences between distributions, suggesting that the category of an influencer does not affect the number of retweets they receive. Additional results are provided in Supplementary Information, including the prevalence of each influencer category across retweet count buckets (Supplementary Fig. S3).

To further explore the role of influencer categories, we test whether the retweet distributions align with a null model in which these labels are assigned randomly. Two randomization strategies are applied:

- Reshuffling: Category labels are randomly reshuffled across influencers, to ensure each category keeps its original size.
- 2. Uniform Sampling: Categories are randomly sampled from a uniform distribution, with each having an equal probability of selection (p = 1/6).

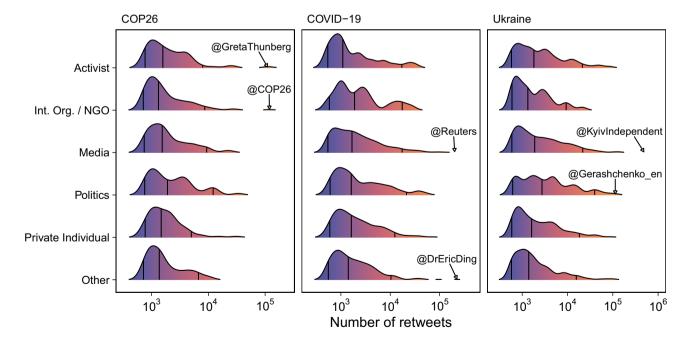


Fig. 1. Distribution of the number of retweets received by each influencer category and for each debate. The vertical lines reported over each distribution indicate, from left to right, the 0.05, 0.5, and 0.95 quantiles. We also report the usernames of the two most retweeted accounts per debate. No category appears to dominate over the others, but rather singular prominent accounts in each debate emerge over the rest. Notable examples are @GretaThunberg in the climate debate and @KyivIndependent in the debate about the Russo-Ukrainian War.

We compare the observed and randomized distributions using two-sided two-sample Wilcoxon rank-sum tests for both strategies. The results show no significant differences, indicating that the observed number of retweets is consistent with random assignment of categories. This reinforces the finding that an actor's category does not appear to affect their ability to attract retweets. Detailed results are provided in Supplementary Table S2, as well as visualizations of the randomized distributions (Supplementary Fig. S4).

Finally, we analyze users' preferences in retweeting influencer categories. To ensure that all users in a dataset might have retweeted each category at least once, we restrict our analysis to users with at least six retweets and compute how frequently they retweet each category. We then employ the Gini Index $(GI)^{70}$ to measure how evenly users distribute their retweets across influencer categories. The GI is a value in [0,1] that quantifies inequality within a distribution: a value of 0 corresponds to minimum inequality (i.e., all categories are retweeted by the user with the same frequency), while a value of 1 corresponds to maximum inequality (i.e., the user only retweets a single category). See Methods for its formal definition.

The resulting distributions of users' GI, displayed in Supplementary Fig. S5, suggest that users show a preference for a limited subset of categories. For COP26, the GI distribution has a median of 0.73 (interquartile range: $Q1=0.60,\ Q3=0.87$); for COVID-19, the median is 0.77 ($Q1=0.65,\ Q3=0.89$); and for Ukraine, the median is 0.69 ($Q1=0.58,\ Q3=0.80$). However, when considering the total number of retweets each category receives, we observe lower levels of inequality across all debates: GI = 0.28 for COP26, GI = 0.59 for COVID-19, and GI = 0.53 for Ukraine.

The difference between the GI estimates at the user and debate level indicates that, while most users tend to favor certain categories over others, the specific subset of preferred categories varies across individuals. Users retweet different influencer categories, but these individual preferences balance out, resulting in smaller overall inequality across categories. This pattern is visible in all three debates, particularly in the case of COP26.

Ideological persistence as a driver of engagement

Our analyses show that considering only influencers' categories is insufficient to explain the drivers of user engagement on social media. As other factors such as ideological alignment⁷¹ have been shown to influence how people interact with social media content, we now investigate the role of ideology in shaping engagement across the key debates analyzed in this study.

the key debates analyzed in this study.

In line with previous research^{71–73}, we estimate the ideological leaning of an individual i using the "latent ideology" method. This technique infers the ideological distribution within each debate by assigning each individual a scalar value x_i representing their ideological position. These values are derived from the users' interactions with influencers, allowing us to capture potential patterns of ideological polarization.

We begin by constructing an interaction matrix where each element corresponds to the number of times a user retweets a specific influencer. Then, the method assigns numerical scores to users so that individuals retweeting similar subsets of influencers receive similar scores. By rescaling these values to the range [-1, +1], we obtain an ideological spectrum, with users at opposing ends interacting with distinct and non-overlapping sets of influencers. A detailed explanation of the technique is provided in Methods.

Following the methodology adopted in previous works⁷¹, we focus on the top 300 most retweeted influencers from each of the three debates. We calculate the latent ideology of all users within each debate and validate the resulting scores to ensure that the distributions represent true ideological spectra (see Supplementary Section 2.1 and Supplementary Fig. S6). Finally, we follow standard procedure and estimate the ideological position of each influencer as the median ideology score of their retweeters⁷¹.

Figure 2 illustrates the latent ideology distributions for users and influencers across the three debates. To ensure consistency, we map the majority group to -1 and the minority group to +1 for all cases. Although the extent of polarization varies between debates, all distributions exhibit a clear bimodal structure, with most users clustered near the spectrum's extremes. This observation indicates the existence of two dominant, ideologically opposing communities within each debate.

To formally evaluate the previous observation, we compute Hartigan's dip statistic to test for unimodality (see Methods). The results are consistent across debates, with D=0.021~(P<0.001) for COP26, D=0.035~(P<0.001) for COVID-19, and D=0.068~(P<0.001) for Ukraine. These values indicate a statistically significant deviation from unimodality, confirming the presence of two well-separated communities in each debate. These ideologically opposing clusters reflect the presence of two polarized communities, where users predominantly interact with influencers aligned with their ideological stance.

Having manually annotated all influencers, we can identify the ideological stances most representative of these communities. Influencers in the Majority predominantly advocate for pro-climate action, pro-vaccination, and pro-aid stances in the COP26, COVID-19, and Ukraine discussions, respectively. Conversely, influencers in the Minority are typically associated with alternative or fringe viewpoints, such as skepticism toward climate change and COVID-19 vaccination. For instance, accounts with large followings such as @JunkScience for COP26, @HighWireTalk for COVID-19, and @georgegalloway for Ukraine, were all assigned to the ideological Minority.

Consistent with previous research on online polarization^{8,71,74,75}, our findings reveal the presence of communities supporting conflicting ideologies. This dynamic fosters user engagement with one of these opposing narratives, resulting in a recurring pattern of polarization across debates.

To evaluate the persistence of ideological segregation between different topics, we quantify the ideological overlap of users across debates. Figure 3 presents the density plot of the joint distribution P(x, y) of users' ideologies, where P(x) and P(y) correspond to the ideology distributions for the debates represented on the x and y axes, respectively. This analysis involves N=43,392 users with an ideology score both in the COP26 and COVID-19 debates, N=91,647 in COVID-19 and Ukraine, and N=28,463 in Ukraine and COP26.

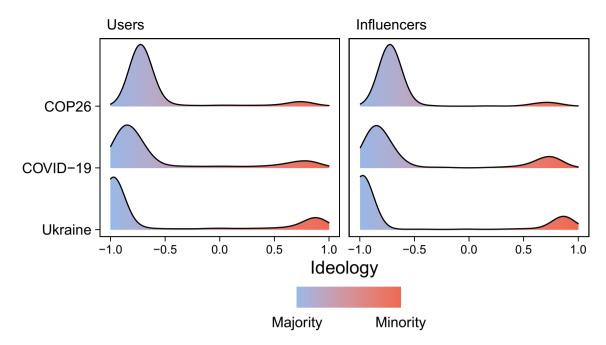


Fig. 2. Ideological spectra of users and influencers in the COP26, COVID-19, and Ukraine debates. The majority group is mapped to -1, whereas the minority group to +1. All distributions exhibit bimodality, indicating the presence in each debate of two main communities of users who retweet distinct sets of influencers.

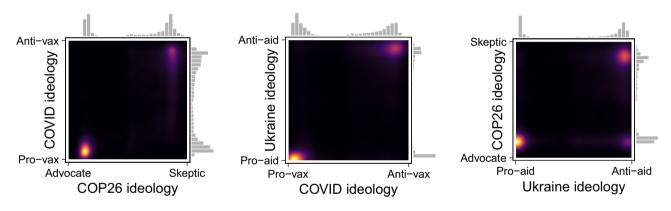


Fig. 3. Joint ideology spectra of users for each pair of debates, with marginal distributions on the top and right sides. The 'Majority' ('Minority') ideology is mapped to the left (right) of the *x*-axis and the bottom (top) of the *y*-axis. Colors indicate the density of users in a region, with lighter colors corresponding to a higher concentration. Most users occupy the regions on the bottom-left and the top-right of the ideology space, indicating a substantial overlap between the Majority and Minority communities across debates.

In the figure, points along the bisector of the first and third quadrants represent users who maintain a consistent ideological stance (Majority or Minority) across both debates. Conversely, points along the bisector of the second and fourth quadrants represent users with misaligned stances between debates (e.g., Majority in one debate and Minority in the other). To account for the size imbalance between the Majority and Minority groups shown in Fig. 2, we generate each panel of Fig. 3 by first downsampling P(x) and then P(y), to ensure that the number of users in the Majority group matches that of the Minority group. The full distributions without downsampling are provided in Supplementary Fig. S7.

Our findings reveal a substantial overlap of ideological communities across debates. This is reflected by the high density of users in the bottom-left (Majority in both) and top-right (Minority in both) corners of the joint distributions. Interestingly, the joint distribution of Ukraine and COP26 ideologies reveals a subset of users who hold a minority, anti-aid stance while simultaneously supporting climate action, such as @ReadeAlexandra and @DeborahDupre. This divergence suggests that the strength of ideological alignment may vary depending on the topics under discussion.

To further assess the persistence of the Majority versus Minority configuration across debates, we calculate the probability of a user maintaining the same ideological stance across different topics. Specifically, we estimate the fraction of users within the Majority or Minority group in one debate who belong to the same group in the other two debates.

The results, presented in Fig. 4, further support the qualitative insights derived from the joint distribution in Fig. 3. The two matrices display the fraction of users who remain in the majority (or minority) group in a debate on the *x*-axis, given that they belong to the majority (or minority) group in the debate specified on the *y*-axis. This asymmetry arises from the different fractions of majority and minority supporters in each debate. For instance, the proportion of COP26 supporters who also endorse COVID-19 vaccines may differ from the proportion of vaccine supporters who also back climate change policies. The matrices reveal that a user's stance in one debate is highly indicative of their stance in others, with the proportion of users maintaining the same ideological leaning equaling or exceeding 90% for most debate pairs. This result is robust to using more restrictive ideology thresholds to label users as 'Majority' or 'Minority' (Supplementary Fig. S8).

Users in the Majority group of one debate are highly likely to belong to the Majority group in another. For instance, 99% of users in the ideological Majority of the COP26 debate also belong to the Majority in the COVID-19 debate. Similar patterns are observed for users in the Minority group, albeit with slightly less consistency. Most users in the Minority of one debate tend to adopt a Minority stance in the others. Notably, the Minority group in the COVID-19 debate exhibits a high level of ideological cohesion, with approximately 90% of these users also holding a Minority position in the COP26 and Ukraine debates. In contrast, the Ukrainian debate stands out as the only case where a Minority stance does not strongly indicate corresponding positions in both other debates.

This discrepancy may be attributed to the distinct characteristics and perceived salience of the debates under study. The COVID-19 pandemic, as a global and prolonged event, likely fostered a cohesive ideological identity among users skeptical of institutional narratives. This shared identity appears to extend across debates, such as COP26, reflecting a broader distrust of mainstream scientific and institutional messaging. Conversely, the Russo-Ukrainian War, while significant, represents a geographically and politically localized event. This localized nature may result in less ideological cohesion among Minority users across debates. Additionally, the Ukraine debate's complexity— encompassing geopolitical perspectives, humanitarian concerns, and diverse narratives—could contribute to weaker alignment with stances on other topics.

On social media, ideological alignment has a greater impact on shaping discussions than the social, political, or professional roles of individuals or organizations. No single category of actors consistently drives the discussion. Instead, our findings highlight the persistence of a Majority versus Minority ideological configuration across debates. Users tend to consistently align with one side of the spectrum, suggesting that the structure of online discourse is less about the hierarchy of different types of actors and more about the alignment of narratives with pre-existing ideological stances.

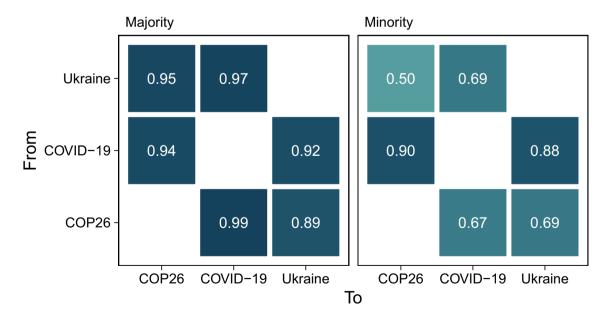


Fig. 4. Conditional probability of user ideology across debates. Each tile reports the fraction of users in a 'From' debate who belong to the same community in a 'To' debate. More than 90% of users with 'Majority' ideology on a debate (left panel) stay in the majority community of the other two. Similarly, but to a lesser extent, users with 'Minority' ideology on a debate (right panel) tend to stay in the minority community of the other two, especially users within the ideological minority of the COVID-19 debate.

Influencer engagement patterns in majority and minority communities

In the previous sections, we show that most influencers receive a similar number of retweets, irrespective of their category. Furthermore, users exhibit consistent behavior by retweeting influencers aligned with the Majority or Minority ideological communities across topics. Now, we aim to investigate further how the two communities differ in terms of the influencers they retweet.

To begin, we analyze whether the number of retweets an influencer receives varies depending on the ideological community they belong to. For this purpose, we employ a bootstrapping procedure to estimate the median number of retweets and retweeters attracted by influencers associated with the Majority and Minority communities. The results, presented in Fig. 5 and detailed in Supplementary Table S3, show that across all three debates, the estimates for the two groups of influencers are comparable. This indicates that individual influencers attract similar numbers of retweets and unique retweeters, regardless of their ideological alignment, and suggests that the ideological community to which an influencer belongs does not substantially impact the engagement they receive.

Next, we investigate which categories are more frequently retweeted by users with a Majority or Minority stance. For this purpose, we construct a network of influencer categories for each debate and ideological stance. In these networks, two categories are connected if a user retweets influencers from both categories at least once, with the weight of each connection representing the number of distinct users who retweeted both. Figure 6 illustrates the networks generated from this procedure.

In general, the resulting networks indicate that users in the Majority and Minority communities tend to retweet influencers from distinct categories. This distinction is particularly clear in the COP26 debate, where 35% of all retweets to influencers by users in the Minority are received by Private Individuals, while the Majority primarily engages with Activists (21%) or Organizations (28%). Additionally, Media accounts receive a higher relative number of retweets from users in the Minority (35%) compared to those in the Majority (23%), reflecting a distinction in how the two ideological communities interact with information sources.

Similar patterns emerge in the Ukraine debate, albeit less pronounced. Media and Private Individuals are the more popular influencer categories for both ideological communities, accounting for 68% and 71% of all retweets to influencers by Majority and Minority users, respectively. In contrast, other categories are more popular in one community than the other, such as Politics (15% of total retweets by the Majority and 5% by the Minority).

Interestingly, users in the Majority and Minority of the COVID-19 debate display similar preferences, as both communities tend to retweet each category with similar frequency relative to one another. This lack of distinction indicates that no influencer category is favored more by one ideological community over the other, as we observe instead in other debates.

These findings further reinforce the conclusion that users' engagement on social media is shaped more by ideological alignment than by influencers' professional or social roles. This dynamic reflects the polarization observed in online discourse, where individuals selectively interact with content that aligns with their preexisting views.

Conclusions

Our study provides new insights into the interplay between user engagement and influencers' roles in major global debates on social media. By analyzing large-scale datasets from the COP26, COVID-19, and Russo-Ukrainian War debates, we demonstrate that user ideological alignment holds greater significance than an influencer's category in shaping online discussions.

In traditional media environments, agenda-setting theories emphasize the hierarchical influence of institutional actors, such as mainstream media and political elites, in prioritizing public issues. Our findings suggest that these dynamics may not fully extend to social media. While certain users achieve prominence within specific topics, their relevance does not appear to be tied to their social or professional role, nor is it consistent across debates. Rather, the engagement they receive may arise from resonance within polarized communities.

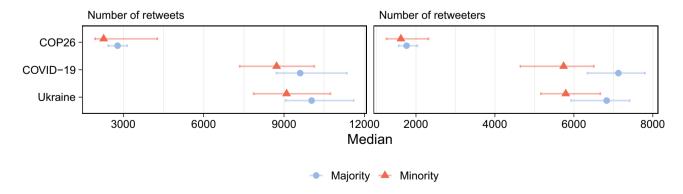


Fig. 5. Median number of retweets and retweeters attracted by influencers of the Majority and Minority communities. The estimates result from a bootstrapping procedure, with error bars corresponding to the 95% bias-corrected and accelerated confidence intervals. Across all debates, the two sets of influencers receive a comparable number of retweets from a similar number of unique users.

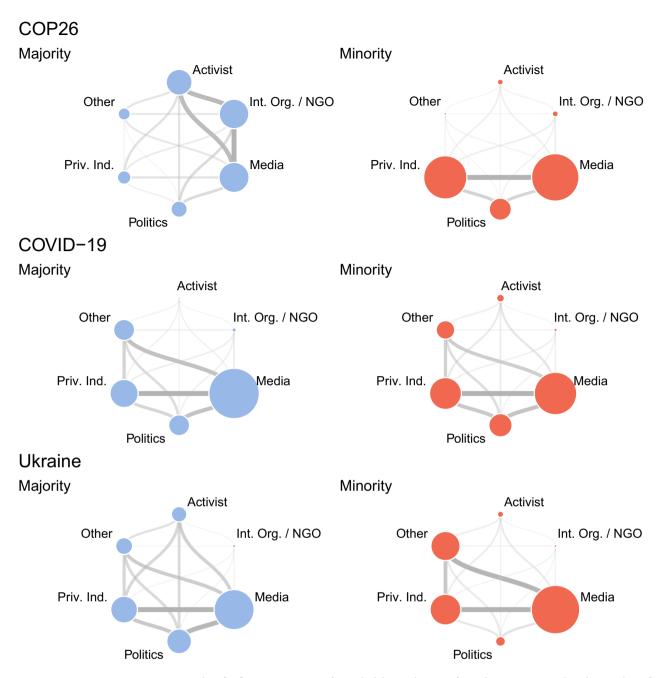


Fig. 6. Networks of influencer categories for each debate. The size of a node is proportional to the number of users who retweeted that category. The width and visibility of an edge between two categories is proportional to the number of users who retweeted both. Only retweets by users in the Majority are considered on the left column, whereas on the right column, only those by users in the Minority. The resulting networks show that the users of the two communities tend to retweet different categories of influencers. This is particularly evident in the case of the COP26 and Ukraine debates, whereas differences are not as apparent in the case of COVID-19.

These findings hint at a change in the traditional top-down agenda setting model, pointing to a growing role for decentralized, bottom-up mechanisms, where users collectively amplify narratives aligned with their ideological preferences. The persistence of a Majority versus Minority ideological configuration in all debates highlights the presence of segregated and ideologically contrasting communities that extend beyond specific topics. This consistent affiliation may further limit exposure to diverse perspectives, thereby reinforcing the effects of echo chambers. Further, despite the difference in the size of their communities, the comparable engagement levels obtained by Majority and Minority influencers suggest that their popularity is not directly related to the ideological stance they assume.

While ideological alignment is a unifying factor, our analysis also reveals topic-specific variations. For example, the Ukraine debate exhibits weaker ideological cohesion compared to COP26 and COVID-19,

particularly among Minority users. This difference likely reflects the varying salience and scope of these topics. The COVID-19 pandemic, as a global and prolonged crisis, appears to have fostered a cohesive ideological identity among users with alternative views, extending across debates such as climate change. In contrast, the Ukraine conflict, characterized by regionally and politically fragmented narratives, contributes to weaker ideological alignment across topics.

Our results also offer some noteworthy applications. For communication professionals, they suggest a need to rethink outreach strategies: rather than targeting audiences solely based on topic relevance or demographic features, it may be more effective to account for users' broader ideological positioning. Tailoring messages that bridge ideological narratives or are endorsed by trusted influencers within specific communities could increase the chance of cross-cluster diffusion. However, the susceptibility of this type of dynamic shall be taken into account also from a policymaking point of view, as it may fuel the cross-spreading of radicalized narratives between different sensitive topics. Additionally, our findings contribute to the understanding of filter bubbles by demonstrating that they are not isolated to specific issues but reflect systemic, cross-topic ideological alignment. This presents challenges for initiatives aimed at promoting exposure to diverse perspectives. Efforts to mitigate filter bubbles might benefit from designing interventions that operate across topics and target underlying ideological structures, rather than focusing narrowly on specific debates.

The decentralized nature of social media discourse fostered a new model of information spreading, challenging the traditional gatekeeping power of dominant actors. This new paradigm presents both opportunities and risks. On one hand, it democratizes the flow of information, amplifying diverse perspectives. On the other hand, it risks exacerbating polarization and misinformation, as users selectively engage with ideologically congruent content across topics. Moreover, as the traditional filters that once ensured a baseline of editorial standards have weakened, questions arise about the quality of content consumed in online environments. In turn, the fragmentation of the information space undermines the conditions for a shared public debate, an essential foundation of any democratic society.

Despite the robustness of our analyses, several limitations must be acknowledged. First, our investigation focuses exclusively on Twitter and may not fully reflect user dynamics on other platforms, as platform-specific biases and algorithmic curation could, in principle, influence the patterns we observe. Additionally, the classification of users into Majority and Minority groups relies on latent ideology modeling, which, while effective, may simplify the multifaceted nature of user beliefs.

At the same time, we only focused on English tweets, which has to be addressed as a constraint on the generalizability of the findings. However, this detail also underscores a solidity in our findings. Despite focusing on a language used predominantly by non-local audiences, we still observe strong patterns of ideological alignment and cross-topic coherence. For instance, regarding the Russo-Ukrainian debate, even within English-speaking communities—often external to the conflict zone—polarized narratives persist and align with ideological stances in other unrelated debates (e.g., climate change and COVID-19). This upshot suggests that the ideological dynamics we identify are not only local or topic-specific but reflect a broader structure of online polarization that transcends geographic and linguistic boundaries. Nevertheless, we acknowledge that expanding future analyses to include multilingual datasets could capture more nuances in local perspectives and offer a more comprehensive view of discourse dynamics within involved regions.

Future research should extend these analyses to explore cross-platform dynamics, temporal shifts in ideological alignment, and the role of emerging fringe platforms in shaping discourse.

Methods

Data

All data used in this work were collected via the former Twitter API for academic research, using search queries specific to each dataset. We intentionally selected three heterogeneous debates—climate change (COP26), the COVID-19 pandemic, and the Russo-Ukrainian War—due to their global relevance and high level of online engagement. Although they vary significantly in nature and content, all three debates are ideologically charged and politically salient. This selection allowed us to test whether the mechanisms we identified, including ideological polarization, are structural rather than topic-specific.

Tweets and user information related to the COP26 debate were collected using the 'cop26' search query, from 1 June 2021 to 14 November 2021. The dataset includes approximately 8 million English tweets by 1 million distinct users.

COVID-19 data were collected from Twitter in the period from 1 January 2020 to 30 April 2021 using the official Twitter API through keyword search. We collected every publicly available tweet that contained one of the following keywords: vaccin*, dose*, pharma*, immun*, no-vax*, novax*, pro-vax*, provax*, antivax*, antivax*. The final dataset results in approximately 35 million English tweets made by 8 million unique users.

Tweets related to the Russo-Ukrainian conflict were gathered by tracking specific keywords and accounts from 22 February 2022 to 17 February 2023. We leveraged a publicly available dataset⁷⁶ and hydrated a random sample comprising 25% of all available tweet IDs (https://github.com/echen102/ukraine-russia). The final hydrated dataset includes approximately 85 million English tweets by 7 million distinct users.

Identifying the influencers in a debate

Account selection

For each dataset, we identify a set of 'influencers' that shape discussions within the context of a debate. We note that, by focusing only on English tweets, the vast majority of these influential accounts will also be English-speaking.

We start by identifying the users among a dataset's top 1% most retweeted accounts. Users whose content is extensively reshared may have a role in driving discussions on social media, as retweeting may increase the

visibility of certain content and, in turn, heighten its perceived salience in the public eye. However, focusing solely on retweets may also yield accounts that have received many retweets, but only on a small number of tweets that may have gone viral. Instead, we are interested in users who have been consistently influential in the debate, by not only receiving many retweets but also producing a sufficiently large quantity of original content. For this reason, we restrict the resulting sets of influencers in each dataset to only include the top 50% producers of original tweets (i.e., tweets that are not replies, quotes, or retweets).

Before any manual removal of accounts that may have been wrongly included (e.g., if their content is unrelated to the dataset's theme) or that could not be classified into a category (see Supplementary Section 1.1), this selection results in 603 influencers for the COP26 debate, 2168 for the COVID-19 debate, and 1385 for the Ukraine debate.

Manual annotation

After identifying a set of influencers for each debate, we manually label each account using ad-hoc questionnaires, in a way similar to⁷¹. This labeling procedure enables us to associate with each user several features that are not easily retrievable or inferable from their tweets or account information.

We assign each account to one of six predefined categories based on their social, political, or professional role, namely 'Activist', 'International Organization/NGO', 'Media', 'Politics', 'Private Individual', and 'Other'. These categories are shared across all debates. Additionally, we label each account with its main ideological stance in the debate, such as their support for climate action in the COP26 debate, vaccination in the COVID-19 debate, and military aid to Ukraine. We note that the influencer set for each debate is annotated independently of the others, meaning an influencer is labeled with respect to a specific debate only if they appear among the influencers of that particular debate.

The full list of questions, in addition to further details about the annotation procedure, is reported in Supplementary Section 1.1.

Latent ideology

The latent ideology technique 72,73 allows for estimating a synthetic opinion distribution using retweet interactions. Following a procedure detailed in prior works 71 , we begin by constructing an interaction matrix M, where each element M_{ij} is equal to the number of times user i has retweeted user j. We exclusively focus on the interactions between users and influencers (as defined in the previous section), with matrix rows corresponding to the former and matrix columns to the latter. Additionally, we ignore all users i such that $\sum_j M_{ij} \leq 1$. Next, we run Correspondence Analysis 77 on matrix M. The method requires the following steps. First, we define a new matrix P obtained from the normalization of M with the total number of retweets

$$P = \frac{M}{\sum_{ij} M_{ij}}$$

and then consider the vector of row sum $\mathbf{r} = P1$ and column sum $\mathbf{c} = \mathbf{1}^T P$. Then, we compute the matrix of standardized residuals of M as

$$S = D_r^{-1/2} (P - \text{rc}) D_c^{-1/2}$$

where $D_r = \operatorname{diag}(\mathbf{r})$ and $D_c = \operatorname{diag}(\mathbf{c})$. Finally, we apply single value decomposition to matrix S as

$$S = U D_{\alpha} V^{T}$$

where $UU^T=VV^T=I$ and D_α is the matrix of singular values of S. User ideologies are estimated as the standard row coordinates $X=D_r^{-1/2}U$. We focus only on the first dimension, corresponding to the largest singular value, and rescale it to [-1,+1]. Like previous works^{71,73}, influencer ideologies are estimated as the median ideology scores of their retweeters.

Hartigan's dip test

Hartigan's dip test is a statistical test designed to assess the unimodality or multimodality of a distribution 78 . Specifically, it tests the null hypothesis of an unimodal empirical distribution (i.e., having a single peak) against the alternative hypothesis that it is multimodal (i.e., having more than one peak). The test computes a statistic D that quantifies the extent of multimodality, with larger values corresponding to larger multimodality. The statistical significance p of the test is used to assess the multimodality of the distribution, with low p-values indicating that the hypothesis of unimodality can be rejected. Notably, this test can be used to assess whether an empirical ideology distribution can be considered bimodal. In this work, we used this test to assess whether the user opinion distributions inferred with the latent ideology method can be considered bimodal (Fig. 2), as distributions displaying two peaks at the opposing ends of a spectrum are characteristic of a polarized debate.

Gini index

The Gini Index (or coefficient) is a measure that quantifies the inequality in the frequency values of a distribution 70 . For a vector x of n frequencies, it is defined as

$$GI = \frac{\sum_{i=1}^{n} \sum_{j=1}^{n} |x_i - x_j|}{2n^2 \overline{x}}$$

where $\overline{x} = \sum_{i=1}^n x_i/n$ is the mean frequency. The measure is defined in [0,1], with $\mathrm{GI} = 0$ corresponding to perfect equality $(x_i = x_j, \ \forall i \neq j)$ and $\mathrm{GI} = 1$ to perfect inequality $(x_i = 1 \ \mathrm{and} \ x_{j \neq i} = 0)$. We use this measure to quantify the inequality of the proportion of retweets received by each category, both at the user level and the debate level.

Data availability

Twitter data are made available under the platform's terms of service. Tweet IDs for COP26 are available on OSF (https://osf.io/nu75j/). Tweet IDs for COVID-19 and Ukraine will be made available upon request to the corresponding author.

Code availability

The R code used to compute the latent ideology scores is available on OSF (https://osf.io/nu75j/). All other analyses were carried out using base R and freely available packages.

Received: 7 January 2025; Accepted: 10 September 2025

Published online: 14 October 2025

References

- 1. Bakshy, E., Rosenn, I., Marlow, C. & Adamic, L. The role of social networks in information diffusion. In Proceedings of the 21st international conference on World Wide Web, 519–528 (2012).
- 2. Tucker, J. et al. Social media, political polarization, and political disinformation: A review of the scientific literature. SSRN Electron. J. (2018).
- 3. Aichner, T., Grünfelder, M., Maurer, O. & Jegeni, D. Twenty-five years of social media: a review of social media applications and definitions from 1994 to 2019. Cyberpsychology, behavior, and social networking 24, 215–222 (2021).
- Di Marco, N., Cinelli, M., Alipour, S. & Quattrociocchi, W. Users volatility on reddit and voat. IEEE Transactions on Comput. Soc. Syst. 11, 1–9 (2024).
- 5. Van Bavel, J. J., Rathje, S., Harris, E., Robertson, C. & Sternisko, A. How social media shapes polarization. *Trends Cogn. Sci.* 25, 913–916 (2021).
- Guess, A. M. et al. How do social media feed algorithms affect attitudes and behavior in an election campaign?. Science 381, 398–404 (2023).
- 7. Flaxman, S., Goel, S. & Rao, J. M. Filter bubbles, echo chambers, and online news consumption. *Public opinion quarterly* 80, 298–320 (2016).
- 8. Cinelli, M., De Francisci Morales, G., Galeazzi, A., Quattrociocchi, W. & Starnini, M. The echo chamber effect on social media. *Proc. Natl. Acad. Sci.* 118, (2021).
- 9. Nyhan, B. et al. Like-minded sources on facebook are prevalent but not polarizing. Nature 620, 137-144 (2023).
- 10. Del Vicario, M. et al. The spreading of misinformation online. Proc. Natl. Acad. Sci. 113, 554-559 (2016).
- 11. Shmargad, Y. & Klar, S. Sorting the news: How ranking by popularity polarizes our politics. *Political Communication* 37, 423–446 (2020).
- 12. González-Bailón, S. et al. Asymmetric ideological segregation in exposure to political news on facebook. *Science* **381**, 392–398 (2023).
- 13. Siegel, A. A. Online hate speech. In Social media and democracy: The state of the field, prospects for reform, chap. 4, 56–88 (Cambridge University Press, 2020).
- 14. Castaño-Pulgarín, S. A., Suárez-Betancur, N., Vega, L. M. T. & López, H. M. H. Internet, social media and online hate speech. systematic review. *Aggress. violent behavior* **58**, 101608 (2021).
- 15. Lupu, Y. et al. Offline events and online hate. PLoS one 18, e0278511 (2023).
- 16. Avalle, M. et al. Persistent interaction patterns across social media platforms and over time. Nature 628, 582-589 (2024).
- 17. Loru, E., Cinelli, M., Tesconi, M. & Quattrociocchi, W. The influence of coordinated behavior on toxicity. *Online Soc. Networks and Media* 43–44, (2024).
- 18. Ji, P. et al. Signal propagation in complex networks. Phys. Reports 1017, 1-96. https://doi.org/10.1016/j.physrep.2023.03.005 (2023).
- Van Aelst, P. et al. Political communication in a high-choice media environment: a challenge for democracy?. Annals Int. Commun. Assoc. 41, 3–27 (2017).
- 20. Anderson, S. P. & De Palma, A. Competition for attention in the information (overload) age. The RAND J. Econ. 43, 1-25 (2012).
- 21. Webster, J. G. The marketplace of attention: How audiences take shape in a digital age (Mit Press, 2014).
- 22. Lorenz-Spreen, P., Mønsted, B. M., Hövel, P. & Lehmann, S. Accelerating dynamics of collective attention. *Nat. communications* 10, 1759 (2019).
- 23. Bruns, A. 3.1. the active audience: Transforming journalism from gatekeeping to gatewatching (2008).
- 24. Livingstone, S. The participation paradigm in audience research. The communication review 16, 21-30 (2013).
- 25. Webster, J. G. & Ksiazek, T. B. The dynamics of audience fragmentation: Public attention in an age of digital media. *J. communication* 62, 39–56 (2012).
- Perloff, R. M. The fifty-year legacy of agenda-setting: Storied past, complex conundrums, future possibilities. Mass Commun. Soc. 25, 469–499 (2022).
- 27. Prior, M. Post-broadcast democracy: How media choice increases inequality in political involvement and polarizes elections (Cambridge University Press, 2007).
- 28. Gallup, K. Indicators of news media trust (2018).
- 29. Nic, N., Fletcher, R., Kalogeropoulos, A., Levy, D. A. & Nielsen, R. K. Reuters institute digital news report 2018. Reuters Institute for the Study of Journalism 39 (2018).
- 30. Sangiorgio, E., Cinelli, M., Cerqueti, R. & Quattrociocchi, W. Followers do not dictate the virality of news outlets on social media. PNAS Nexus 3, pgae257 (2024).
- 31. Chadwick, A. The hybrid media system: Politics and power (Oxford University Press, 2017).
- 32. Bawden, D. & Robinson, L. *Information overload: An overview* (In Oxford Encyclopedia of Political Decision Making (Oxford University Press, 2020).
- 33. Hargittai, E., Neuman, W. R. & Curry, O. Taming the information tide: Perceptions of information overload in the american home. *The Inf. Soc.* 28, 161–173 (2012).
- 34. Schmidt, A. L. et al. Anatomy of news consumption on facebook. Proc. Natl. Acad. Sci. 114, 3035–3039 (2017).
- 35. Cinelli, M. et al. Selective exposure shapes the facebook news diet. PloS one 15, e0229129 (2020).
- 36. Bessi, A. et al. Science vs conspiracy: Collective narratives in the age of misinformation. PloS one 10, e0118093 (2015).

- 37. Zollo, F. et al. Debunking in a world of tribes. PloS one 12, e0181821 (2017).
- 38. Bennett, W. L. & Iyengar, S. A new era of minimal effects? the changing foundations of political communication. J. communication **58**, 707–731 (2008).
- 39. Bail, C. A. et al. Exposure to opposing views on social media can increase political polarization. Proc. Natl. Acad. Sci. 115, 9216-9221 (2018)
- 40. Deuze, M. The changing context of news work: Liquid journalism for a monitorial citizenry. Int. journal Commun. 2, 18 (2008).
- 41. Singer, J. B. User-generated visibility: Secondary gatekeeping in a shared media space. New media & society 16, 55-73 (2014).
- 42. McCombs, Maxwell E. & D. L. S. & Weaver, D. H., New directions in agenda-setting theory and research. Mass Commun. Soc. 17, 781-802 (2014).
- 43. Barberá, P. et al. Who leads? who follows? measuring issue attention and agenda setting by legislators and the mass public using social media data. Am. Polit. Sci. Rev. 113, 883-901 (2019).
- 44. Galeazzi, A., Peruzzi, A., Brugnoli, E., Delmastro, M. & Zollo, F. Unveiling the hidden agenda: Biases in news reporting and consumption. PNAS Nexus 3, pgae474 (2024).
- 45. Boyd, D. Social network sites as networked publics: Affordances, dynamics, and implications. In A networked self, 47-66 (Routledge, 2010).
- 46. Shoemaker, P. Gatekeeping. newburry park (1991).
- 47. Shoemaker, P. J., Eichholz, M., Kim, E. & Wrigley, B. Individual and routine forces in gatekeeping. Journalism & Mass Commun. Q. 78, 233–246. https://doi.org/10.1177/107769900107800202 (2001).
- 48. Bowman, S. & Willis, C. We media. How audiences are shaping the future of news and information 66, 13-20 (2003).
- 49. Gillmor, D. We the media: Grassroots journalism by the people, for the people ("O'Reilly Media, Inc.", 2006).
- 50. Bro, P. & Wallberg, F. Digital gatekeeping: News media versus social media. Digit. Journalism 2, 446-454. https://doi.org/10.1080/ 21670811.2014.895507 (2014).
- 51. Pantic, M. & Ziek, P. Gatekeeping in a digital media habitat: The role of secondary gatekeepers. Electron. News 19, 94-110. https:/ /doi.org/10.1177/19312431241281687 (2024).
- 52. Hoskins, A. & O'Loughlin, B. Remediating jihad for western news audiences: The renewal of gatekeeping?. Journalism 12, 199-216. https://doi.org/10.1177/1464884910388592 (2011).
- 53. Meraz, S. & Papacharissi, Z. Networked gatekeeping and networked framing on #egypt. The Int. J. Press. 18, 138-166. https://doi.o rg/10.1177/1940161212474472 (2013).
- 54. McCombs, M. E. & Shaw, D. L. The agenda-setting function of mass media. The Public Opin. Q. 36, 176-187 (1972).
- 55. Shaw, D. L., McCombs, M., Weaver, D. H. & Hamm, B. J. Individuals, groups, and agenda melding: A theory of social dissonance. Int. J. Public Opin. Res. 11, 2-24 (1999).
- 56. Shaw, D. L., Minooie, M., Aikat, D. & Vargo, C. J. Agendamelding: News, social media, audiences, and civic community (International Academic Publishers, 2019).
- 57. McCombs, M. E., Shaw, D. L. & Weaver, D. H. New directions in agenda-setting theory and research. In Advances in foundational mass communication theories, 131-152 (Routledge, 2018).
- 58. Roberts, M., Wanta, W. & Dzwo, T.-H. Agenda setting and issue salience online. Commun. research 29, 452-465 (2002).
- 59. Wallsten, K. Agenda setting and the blogosphere: An analysis of the relationship between mainstream media and political blogs. Rev. policy research 24, 567-587 (2007).
- 60. Meraz, S. Using time series analysis to measure intermedia agenda-setting influence in traditional media and political blog networks. Journalism & mass communication quarterly 88, 176-194 (2011).
- 61. Johnson, T. J. Agenda setting in a 2.0 world: New agendas in communication (Routledge, 2013).
- 62. Agenda-setting theory meets big data. Russell Neuman, W., Guggenheim, L., Mo Jang, S. a. & Bae, S. Y. The dynamics of public attention. J. communication 64, 193-214 (2014).
- 63. Feezell, J. T. Agenda setting through social media: The importance of incidental news exposure and social filtering in the digital era. Polit. research quarterly 71, 482-494 (2018).
- 64. Ceron, A. Twitter and the traditional media: Who is the real agenda setter? In APSA 2014 Annual Meeting Paper (2014).
- 65. Su, Y. & Borah, P. Who is the agenda setter? examining the intermedia agenda-setting effect between twitter and newspapers. J. Inf. Technol. & Polit. 16, 236-249 (2019).
- 66. Rogstad, I. Is twitter just rehashing? intermedia agenda setting between twitter and mainstream media. J. Inf. Technol. & Polit. 13, 142-158 (2016).
- 67. Towner, T. L. & Muñoz, C. L. Tell me an instagram story: Ephemeral communication and the 2018 gubernatorial elections. Soc. Sci. Comput. Rev. 42, 1288-1312 (2024).
- 68. Kim, R. M. & Anderson, A. The agenda-setting function of social media. In Proceedings of the ACM on Web Conference 2025, 601-613 (2025).
- 69. Yi, H. & Wang, Y. Who is affecting who: the new changes of personal influence in the social media era. Front. Psychol. 13, (2022).
- 70. Gini, C. Measurement of inequality of incomes. The economic journal 31, 124-125 (1921).
- 71. Falkenberg, M. et al. Growing polarization around climate change on social media. Nat. Clim. Chang. 12, 1114-1121 (2022).
- 72. Barberá, P., Jost, J. T., Nagler, J., Tucker, J. A. & Bonneau, R. Tweeting from left to right: Is online political communication more than an echo chamber?. Psychol. Sci. 26, 1531-1542 (2015).
- 73. Flamino, J. et al. Political polarization of news media and influencers on twitter in the 2016 and 2020 us presidential elections. Nat. Hum. Behav. 7, 904-916 (2023).
- 74. Garimella, K., Morales, G. D. F., Gionis, A. & Mathioudakis, M. Quantifying controversy on social media. Trans. Soc. Comput. 1 (2018)
- 75. Cota, W., Ferreira, S. C., Pastor-Satorras, R. & Starnini, M. Quantifying echo chamber effects in information spreading over political communication networks. EPJ Data Sci. 8, 1-13 (2019).
- 76. Chen, E. & Ferrara, E. Tweets in time of conflict: A public dataset tracking the twitter discourse on the war between ukraine and russia. Proc. Int. AAAI Conf. on Web Soc. Media 17, 1006-1013 (2023).
- 77. Benzécri, J. & Bellier, L. L'analyse des données: Benzécri, J.-P. et al. L'analyse des correspondances. L'analyse des données: leçons sur l'analyse factorielle et la reconnaissance des formes, et travaux du Laboratoire de statistique de l'Université de Paris VI (Dunod,
- 78. Hartigan, J. A. & Hartigan, P. M. The dip test of unimodality. The Annals Stat. 13 (1985).

| https://doi.org/10.1038/s41598-025-19776-z

Acknowledgements

The work is supported by IRIS Infodemic Coalition (UK government, grant no. SCH-00001-3391); SERICS (PE00000014); PRIN 2022 "MUSMA" (CUP G53D23002930006) funded by EU - Next-Generation EU - M4 C2 II.1; project CRESP from the Italian Ministry of Health under the program CCM 2022; project SEED N. SP122184858BEDB3. A.G. gratefully acknowledges CY4GATE for the financial support.

Author contributions

A.G. collected the data; E.L., A.G., N.D.M., A.Ba., W.Q. conceived the experiments; E.L. conducted the experi-

ments; E.L., A.G., E.S., N.D.M., M.C., W.Q. analyzed the results; All authors wrote and reviewed the manuscript.

Declarations

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at https://doi.org/1 0.1038/s41598-025-19776-z.

Correspondence and requests for materials should be addressed to E.L. or W.Q.

Reprints and permissions information is available at www.nature.com/reprints.

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

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, 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 you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. 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-nc-nd/4.0/.

© The Author(s) 2025