



OPEN Unveiling the drivers of active participation in social media discourse

Anees Baqir^{1,5,9}, Yijing Chen², Fernando Diaz-Diaz³, Sercan Kiyak⁴, Thomas Louf^{3,5}, Virginia Morini⁶, Valentina Pansanella¹⁰, Maddalena Torricelli⁷ & Alessandro Galeazzi^{1,8}✉

The emergence of new public forums in the form of online social media has introduced unprecedented challenges to public discourse, including polarization, misinformation, and the rise of echo chambers. Existing research has extensively examined these topics by focusing on the active actions performed by users, without accounting for the share of individuals who consume content without actively interacting with it. In contrast, this study incorporates passive consumption data to investigate the prevalence of active participation in online discourse. We introduce a metric to quantify the share of active engagement and analyze over 17 million pieces of content linked to a polarized Twitter debate to understand its relationship with several features of online environments, such as echo chambers, coordinated behavior, political bias, and source reliability. Our findings reveal a significant proportion of users who consume content without active interactions, underscoring the importance of considering also passive consumption proxies in the analysis of online debates. Furthermore, we found that increased active participation is primarily correlated with the presence of multimedia content and unreliable news sources, rather than with the ideological stance of the content producer, suggesting that active engagement is independent of echo chambers. Our work highlights the significance of passive consumption proxies for quantifying active engagement, which influences platform feed algorithms and, consequently, the development of online discussions. Moreover, it highlights the factors that may encourage active participation, which can be utilized to design more effective communication campaigns.

The advent of the digital age has opened up the possibility of fast and direct communication among members of our society. While these technological advancements have expanded public access to information, they may have also fostered several pervasive issues, such as polarization, misinformation, and the emergence of phenomena such as filter bubbles and echo chambers, that could all influence public opinion and negatively impact society^{1–4}. While some signs of this societal fragmentation could already be observed during the twentieth century², the ideological divide between opposite factions has been increasing in the recent past^{3,5,6}.

This radicalization in opinions has been shown to be a clear obstacle to dialogue, consensus, and policymaking^{1,2}. Given the potential impact of such divides on our society, scholars have intensively studied this topic, considering various aspects and different definitions of polarization⁷. To reveal ideological divides in online environments, many social media studies have relied on users' interactions, such as follower-following relationships, linking, or re-sharing^{3,5,8–10}. In these studies, the concept of polarization mainly aligns with the notion of 'interactional' or 'structural' polarization¹¹, which refers to the segregation of users into ideologically distinct groups based on their interactions^{12,13}. Polarized debates are also a fertile environment for the spread of misinformation that may harm society at different levels¹⁴. Falsehoods and unsubstantiated claims have been shown to spread widely in social media^{15–17} and may erode trust in reliable sources¹⁸. Such spreading can be facilitated by the emergence of phenomena such as echo chambers¹⁰ and filter bubbles¹⁹. These two concepts, which share some common ground, refer to the information environment that users experience online and how it could affect the diversity

¹Department of Environmental Sciences, Informatics and Statistics, Ca' Foscari University of Venice, Venice, Italy. ²Department of Network and Data Science, Central European University, Wien, Austria. ³Institute for Cross-Disciplinary Physics and Complex Systems, IFISC (UIB-CSIC), Palma, Spain. ⁴Institute of Media Studies, KU Leuven, Leuven, Belgium. ⁵Fondazione Bruno Kessler, Trento, Italy. ⁶Department of Computer Science, University of Pisa, Pisa, Italy. ⁷Department of Mathematics, City University of London, London, UK. ⁸Department of Mathematics, University of Padova, Padua, Italy. ⁹Northeastern University, London, UK. ¹⁰Istituto di scienza e tecnologia dell'informazione "A. Faedo", Consiglio Nazionale delle Ricerche - CNR, Pisa, Italy. ✉email: alessandro.galeazzi@unipd.it

of views they are exposed to²⁰. However, while the term ‘filter bubbles’ mainly refers to the reduction in the ideological diversity of news a user is exposed to due to the personalization of recommendation algorithms²¹, echo chambers also consider the interactions among users¹⁰. Inside echo chambers, users primarily interact with others who share similar beliefs, which may reinforce their views and create segregated environments where alternative viewpoints and dissenting voices are frequently silenced or ignored²². For these reasons, echo chambers may also exacerbate polarization and favor the spread of misinformation²³, as they trap individuals within their ideological bubbles and limit their exposure to diverse perspectives²⁴.

Although some literature questions the existence²⁵ and effects—especially in terms of exposure to diverse views²⁶—of such structures^{27,28}, the majority of research has found evidence of echo chambers, with only works based on self-reported data finding no proof supporting this hypothesis²⁹. Indeed, echo chambers have been identified in various online domains, including blogs³⁰, forums³¹, and prominent social networks like Facebook and Twitter^{10,32–34}. Over the last decade, there has been a surge of academic research to understand the underlying mechanisms and real extent of echo chambers. Scholars have dedicated substantial efforts to characterize them within online social networks^{10,24,34–37}, and developed indices to gauge their presence and strength^{38–40}. Furthermore, various models have been proposed to elucidate the mechanisms driving the emergence of echo chambers^{39,41,42}. While these models consistently emphasize the role of homophily as the primary catalyst for the echo chamber effect, a diverse range of contributing factors has also been proposed. Such factors include limited attention spans²⁴, selective exposure⁴³, confirmation bias⁴⁴, the silencing effect²², and the role of feed algorithms^{10,45}.

In addition to the interaction structure, other factors such as the platform’s suggestion algorithms also play a crucial role in shaping online debates. These algorithms can either boost or limit the visibility of certain content or narratives, thereby further influencing the spread of information and the diversity of opinions users encounter^{46,47}. Moreover, also bots and coordinated accounts, which are designed to inflate engagement to increase content visibility and promote specific narratives, can influence online discussion⁴⁸. Some studies highlighted that bot-driven coordination can distort political discourse by fostering echo chambers and spreading misinformation⁴⁹, and despite efforts by platforms to curb these activities, coordinated accounts and bots may still hold a significant presence in online debates^{50,51}.

The empirical works mentioned above are largely based on the traces of active actions left by users to interact with the content they were shown. For instance, on Twitter(X), such active actions include liking a tweet, replying to it, or reposting it—also known as retweeting. Yet, users may also consume content without performing any active actions—a phenomenon known as *lurking*. Although precise measurements of lurkers’ prominence are limited, estimates suggest they constitute the majority of social network users, ranging from 75 to 90% of those interested in debates^{52,53}. Researchers have examined lurking behavior by inferring users’ interests⁵⁴, ranking their significance⁵⁵, analyzing their behavior across contexts and platforms⁵⁶, and proposing strategies to encourage participation⁵⁷. However, the definition of lurkers varies by context and timeframe, as users may lurk on some topics while actively engaging in others, or shift between passive and active participation over time.

Further, several definitions of active and passive user consumption have been proposed, and the impact of active participation on public opinion and debates has been extensively studied, with scholars identifying varying degrees of active consumption depending on the perspective from which the issue is analyzed⁵⁸. Understanding the relationship between audience activity and the information they are exposed to is crucial for various reasons, including quantifying media influence, comprehending the interaction dynamics between users and information sources, and developing more effective communication strategies^{59–62}.

Nowadays, social media plays a central role in the current information ecosystem⁴⁸, and researchers have analyzed various online debates, showing that certain types of content and accounts—such as hyperactive users⁶³, bots and coordinated behavior⁶⁴, and public figures⁶¹—can attract more engagement, potentially playing a pivotal role in shaping online discussion. However, many of these studies have focused on engagement dynamics by examining only active actions such as shares, likes, and comments, without considering users who view the content without taking any action. As a result, it remains unclear which factors encourage active participation among users once they are exposed to content.

In our work, we adopt a complementary approach and include the presence of passive users to explore the factors that may encourage active content consumption—specifically, those that increase the proportion of users who take active actions when exposed to content. Additionally, we examine the relationship between active participation and various aspects of the debate, such as echo chambers, ideological bias, and source credibility. By integrating both active actions and passive consumption metrics, our work analyzes audience activity in Twitter discussions, highlighting the factors that correlate with increased active participation.

In particular, by using the recently introduced metric called impressions count to quantify passive consumption our work investigates whether the fraction of users who actively engage with content—meaning they leave explicit feedback actions—depends on user and content characteristics such as the ideological leaning or popularity of the author, the presence of multimedia features, the political bias, or the reliability of the mentioned source. Notably, an impression represents content appearing on a user’s screen, reflecting the frequency of visual engagement. While the impressions count can be used to estimate the proportion of users who leave explicit feedback actions after seeing the content, it does not necessarily reflect the cognitive effort a user spends in processing the presented content.

We introduce a metric based on impression count, called *Active Engagement* (AE), which quantifies the fraction of active user engagement, and analyze 17M+ tweets about a controversial topic, whether to provide military support to Ukraine in the ongoing Russo-Ukrainian war. We use this metric to estimate the prevalence of active consumption, intended here as the presence of explicit feedback actions in the consumption process, generated by different types of content.

The results indicate that heightened active participation is predominantly associated with the presence of unreliable and extremely biased news rather than with the popularity or ideological stance of the content producer. This suggests that the level of active participation does not depend on the echo chamber where the content is generated. Conversely, our results imply that different communities experience similar levels of active consumption.

Our work reveals an absence of a relationship between active engagement in online debates and echo chambers, while highlighting the factors that drive users' active participation. Understanding these relationships is crucial for uncovering, on one hand, the interplay between active consumption and feed algorithms that determine content visibility and contribute to shaping online debates, and on the other hand, for designing communication strategies that are more effective in capturing users' active participation and, thereby, fostering message circulation.

Results

Users' interaction structure

The ongoing debate on whether countries should provide military assistance to Ukraine during its conflict with Russia has generated significant attention from influential figures such as politicians, journalists, and committed citizens. The emergence of echo chambers, environments in which users predominantly interact with like-minded peers and reinforce their beliefs, is a common phenomenon observed in such controversial debates within social networks^{5,33}. Examining the presence of echo chambers around this polarizing topic is our first step to analyze the dependence of active consumption on ideological stance. Hence, we estimated users' stances by Latent Ideology Estimation (see "Methods" section). After selecting a set of influential accounts representing the two opposing factions, referred to as influencers, this procedure infers users' ideology scores based on the influencers they have retweeted. Users who have retweeted similar sets of influencers will share similar ideology values, while those who have retweeted disjoint sets will be placed farther apart.

Panel a of Fig. 1 shows the results of this analysis, with the distribution of both users' and influencers' ideology scores on the top and the retweeters distribution for the top 20 influencers on the bottom (see <https://osf.io/5m3vr/> for the figure with all the influencers). We notice a clear division between users supporting military aid to Ukraine and those opposing it, with the first group's scores tending towards -1 , and the second's to close to $+1$. This bimodal distribution of opinions depicts a polarized discussion where most users unambiguously belong to one camp of the debate. On the left side, we find accounts that explicitly expressed support for Ukraine, such as "@NATO", "@SecBlinken", and the Ukrainian president "@ZelenskyyUa", while on the right side, accounts are associated with the Russian government or have a clear pro-Russia stance. This result allows us to categorize users into two groups based on their estimated ideology score: supporters of Ukraine (ideology score ≤ 0) or opponents (ideology score > 0). To better visualize the presence of echo chambers, panel b of Fig. 1 illustrates the retweet network for users and influencers with a degree higher than 100, color-coded according

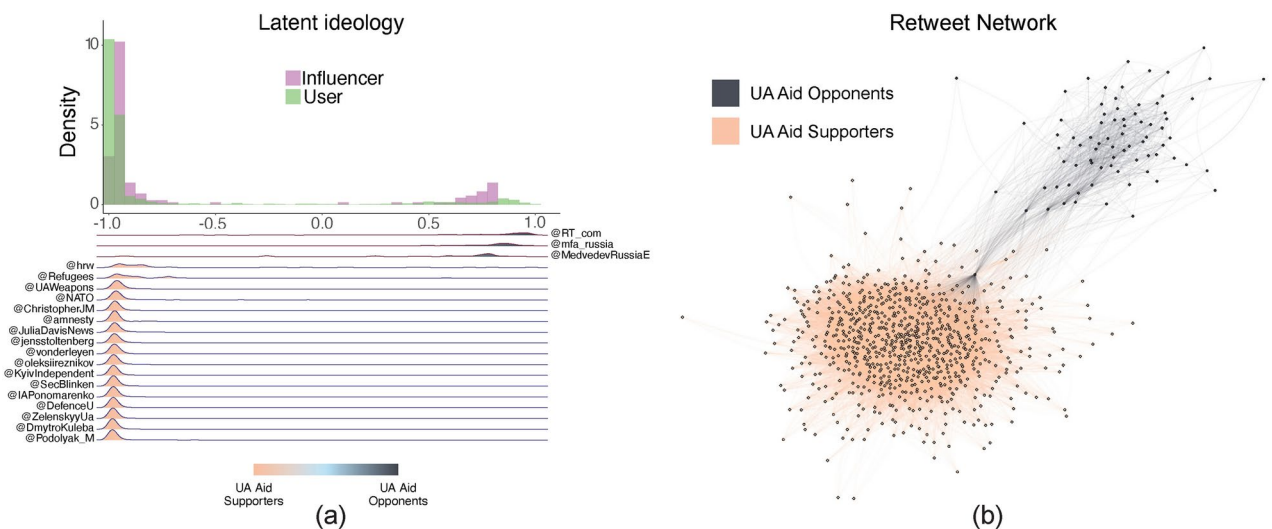


Fig. 1. Latent ideology distribution and retweet network of users and influencers. **(a)** Histogram of users' and influencers' ideology score (top) and retweeters' ideology score distributions for the top 20 influencers (bottom), obtained with the latent ideology estimation algorithm. Negative values represent pro-military aid alignment, while positive values correspond to military aid opponents. Bar colors in the top chart represent the density of influencers (violet) and users (green). The area below retweeter distributions in the bottom chart is shaded in salmon if an influencer is a supporter of Ukrainian aid, and in black if the influencer is against providing weapons to Ukraine. **(b)** Influencers' and users' retweet networks for nodes with a degree greater than 100. Edges and nodes are colored based on nodes' latent ideology values. The two panels reveal the polarized nature of the discussion and the existence of two echo chambers, each endorsing opposing views on Ukraine aid.

to their ideology score. The division between supporters (orange) and opponents (black) of military aid is clear, highlighting the high level of homophilic interaction and underscoring the prevalence of echo chambers (as defined in Cinelli et al.¹⁰) in this debate (see also Fig. S2 in SM). Noticeably, the two groups differ in size, with supporters being nearly 8 times more numerous than opponents.

Unveiling the drivers of active engagement

Having identified the presence of two distinct echo chambers, we then focus on measuring the proportion of the active audience in relation to both the popularity and ideological stance of the users who generated the content, as well as the various characteristics of the content itself. To gauge the proportion of active consumers, we introduce the *Active Engagement* metric. Active Engagement (AE) is calculated as the ratio between the count of one active action, such as the number of likes, retweets, quotes, or replies, and the number of times a piece of content has been viewed, quantified by the “impression count” metric provided by Twitter (see “Methods” section). Notably, impressions quantify appearances, not unique viewers, and hence can be used to estimate the share of passive consumers of content, but not their exact number.

Users’ popularity and ideological stance

We initially explore the correlation between active consumption and content producers’ popularity by comparing the number of followers of the producer with the average Active Engagement (AE) score of their tweets. As shown in Fig. 2a, the majority of users encountering a tweet do not actively interact with it: the average Active Engagement for most users is at least an order of magnitude lower than 1 for all the actions. Further, the graph suggests the dependence of the AE on the type of action, while producers’ popularity seems to have a milder importance. To better quantify the interplay between users’ popularity and AE, we computed the average AE and Pearson’s correlation between content producers’ followers and average Active Engagement for each action (see Table S1 in SM). Quotes exhibit the lowest level of average AE (0.06%), whereas replies and retweets showcase higher and comparable values (0.24% and 0.29% respectively). Likes display the highest AE value (1.1%), suggesting that actions requiring written feedback from users are generally associated with a lower level of active engagement. Notably, we observe a weak to moderate inverse correlation (minimum Pearson’s r coefficient: -0.2250 , maximum: -0.5690 , see Table S1) between the number of followers a user has and her average AE. This suggests that higher popularity does not promise a greater ability to attract active engagement; instead, very popular users seem to experience a slightly higher level of passive consumption.

Next, we shift our focus to examining the relationship between users’ ideological stance and active consumption, exploring the connection between the level of Active Engagement and the echo chamber in which the content is generated. Thus, we utilize the inferred ideology of individual users to compare users’ average AE between the two opposing groups, i.e., Ukraine aid supporters and opponents. Figure 2b shows the boxplots corresponding to the distributions of AE for content generated by supporters and opponents for each action type.

Across the four actions, there are minor differences in the AE distributions between the two communities. Opponents’ retweets, replies, and quotes show a slightly higher median AE, and supporters exhibit wider distributions with longer lower tails. The distributions of likes appear to be very similar, with supporters showing a longer tail of potential outliers. Apart from these minor variations, we do not observe major differences,

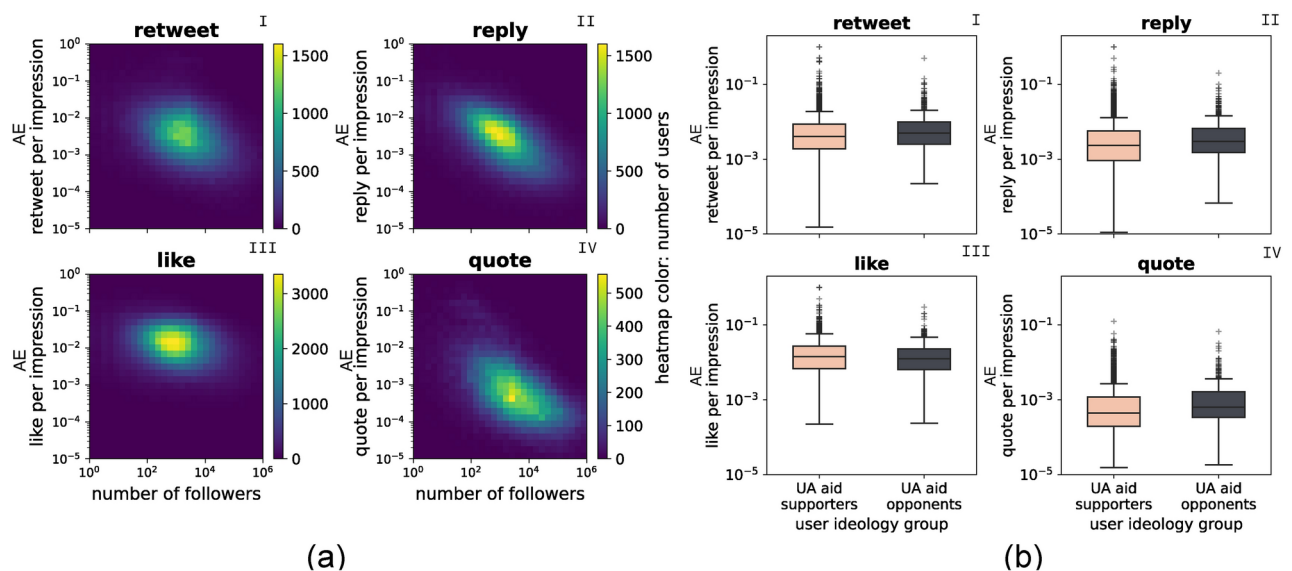


Fig. 2. User-level active consumption for the four active actions. **(a)** Joint distribution density of the number of followers and the active engagement with respect to retweets (panel I), replies (panel II), likes (panel III), and quotes (panel IV). **(b)** Boxplots of the active engagement for the same actions as in panel (a), grouped based on users’ ideologies into UA Supporters and Opponents.

suggesting that despite endorsing opposing ideologies, the two echo chambers do not exhibit pronounced disparities in terms of active consumption levels.

Multimedia content and news sources

In addition to the characteristics of the content producers, some tweet features may also be associated with increased active consumption levels. Specifically, we explore whether the political leaning or reliability of the sources linked in the tweets, along with the presence of multimedia features, correlates with active participation.

Indeed, some studies have found that audiovisual content garners more attention compared to mere text⁶⁵. Hence, we investigate the AE distribution for tweets containing or not containing media (i.e., videos, photos, and GIFs) by relying on the tweet `media_type` field available through the API⁶⁶. Firstly, we observe that the majority of tweets (63.85%) does not contain any media (see Table S2 in SM), and thus the tweets with attached media make up for 36.15% of the original dataset, with a strong prevalence of photos as attached media (30.96%). This distribution highlights that the discussion of UA-aid on Twitter is mainly afforded by text-based content. However, as depicted in Fig. 3a, we find that tweets with media exhibit increased AE compared to those without, although replies are an exception. Specifically, considering tweets incorporating a media, we observe distinct AE distributions across different actions: likes feature a left-skewed distribution with a peak at higher AE values, while quotes exhibit a right-skewed distribution peaking at lower AE values. Replies and retweets, on the other hand, display a more symmetrical distribution. Again, this suggests a dependence of AE on the type of action being considered. Focusing solely on tweets containing visual content (i.e., a video, photo, or animated GIF), the results in Fig. 3b reveal a comparable distribution for photos and videos, with photos generating slightly more active engagement in terms of likes. Notably, GIFs display slightly different patterns with a higher AE in the case of replies and a wider range of AE values in the case of quotes. While the inclusion of any multimedia content increases AE, fewer distinctions are observed across the various types of visual content.

Next, we examine the relationship between the political leaning and reliability of news sources and Active Engagement. Indeed, researchers found that politically loaded and controversial content attracts more attention⁶⁷. Here, we study whether the presence of such content is also correlated with a higher share of active consumers. Employing a source-based approach, we categorize the URLs contained in the tweets based on the domain classification obtained from Media Bias Fact Check (MBFC)⁶⁸ and retrieve labels both for political leaning and reporting reliability. Tweets linked to any news source listed on MBFC can be assigned a political leaning ranging from extreme left to extreme right, and a reliability label from very low to very high (see “Methods” section).

Figure 4a, shows the distribution of AE with respect to the political leaning of the sources in the tweets for the four actions. We notice that highly biased categories, such as extreme right, right, and left, tend to have a higher median engagement and a wider distribution with respect to right center, center, and left center. The extreme left seems to be an exceptional case with a very narrow distribution, but the volume of such tweets is quite low. Yet, these results suggest that heavily biased news can attract more active engagement compared to moderate sources.

Besides political leaning, we also examine the relationship between sources’ reliability and AE. Given previous findings of an imbalance between the consumption of reliable and questionable information^{14–17}, and recognizing the fact that various news sources may have significantly different outreach depending on their popularity, we also consider the number of unique shares as a proxy to estimate sources’ audience size. Figure 4b displays the Active Engagement (AE) plotted against the number of unique shares for news sources, color-coded based on their reliability. Noticeably, AE is not heavily dependent on the popularity (i.e., the number of sharers) of the news sources linked in the tweet, a pattern consistent with our observations in the previous section’s user-level analysis. Moreover, the order of magnitude of AE varies across different types of actions on tweets sharing

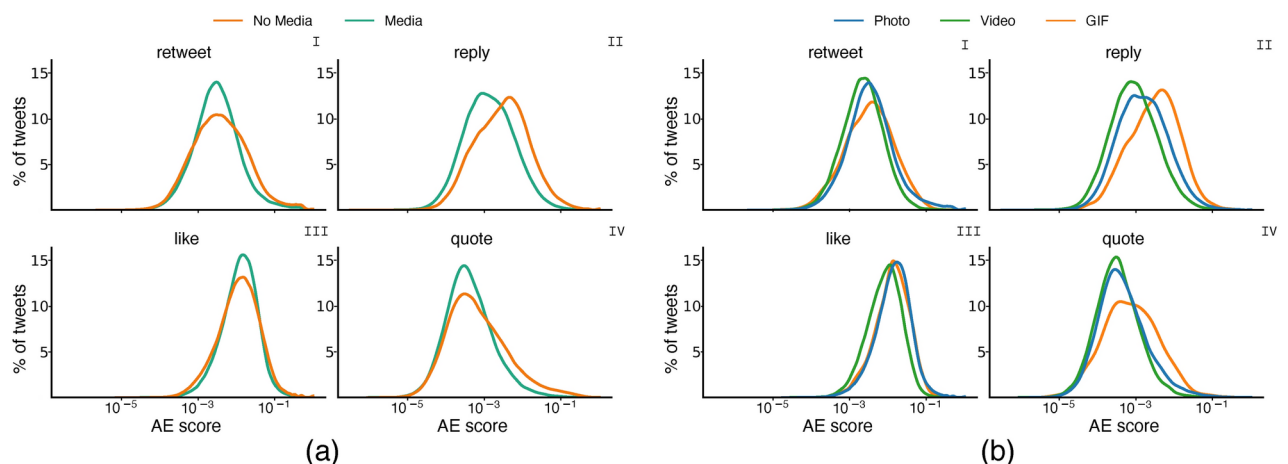


Fig. 3. Active engagement level vs content type. **(a)** Subplots show, for each action, the distribution of active engagement received by tweets having/ not having attached a media. **(b)** Subplots display, for each action, the distribution of active engagement received by tweets having attached different types of media, i.e., photo, video, GIF.

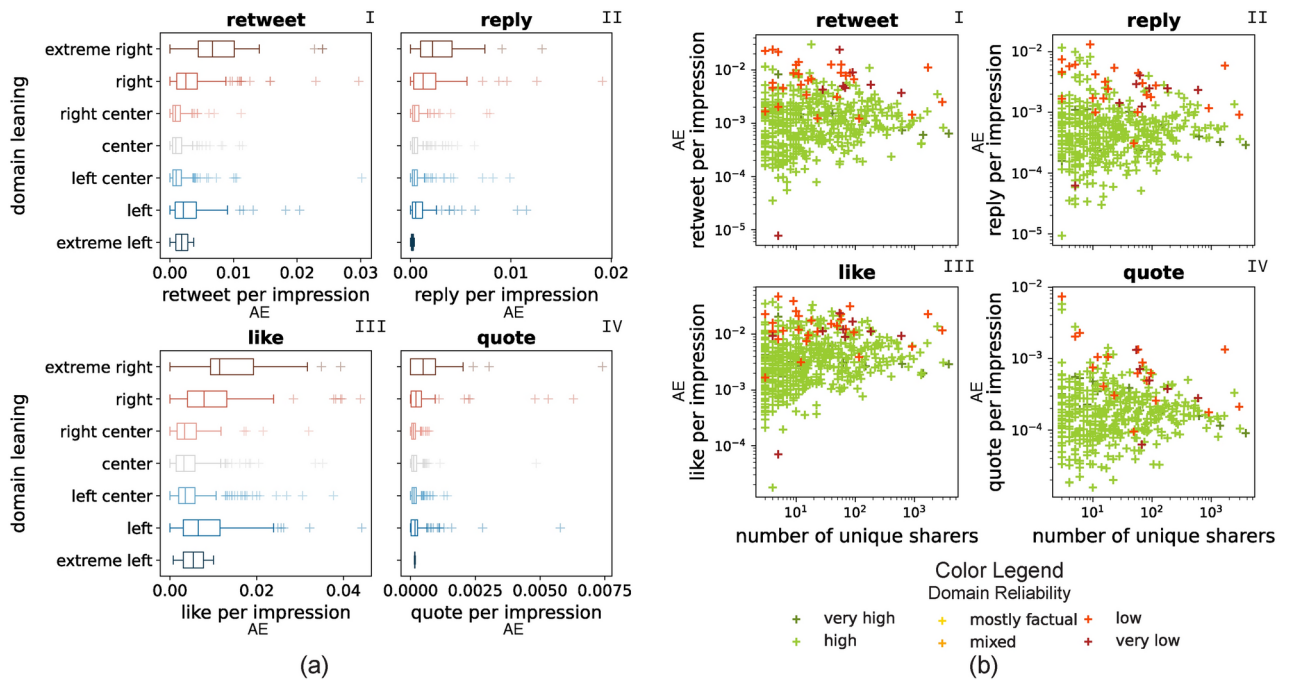


Fig. 4. Influence of news sources' political leaning and reliability on active engagement. **(a)** Boxplot distribution of active engagement with respect to sources' political leaning for each action. **(b)** Active engagement vs the number of unique sharers for news sources, colored by reliability.

domains: the majority of likes maintain a relatively higher AE level around 10^{-3} – 10^{-2} , while the AE of retweets and replies concentrates in a lower range of 10^{-4} – 10^{-3} , with the AE of quotes being the lowest around 10^{-4} . Nevertheless, the impact of low and very low reliability news sources on AE emerges clearly. The two lowest reliability categories appear to be distributed at the top of all plots, implying that tweets linked to these types of sources can attract a higher level of active engagement than those linked to more reliable ones. Among all the users and content characteristics investigated, the presence of low-credibility sources seems to be the most characterizing factor associated with increased levels of AE.

Since previous literature^{3,5} found an overlap between users' ideology and news source usage, as a final analysis we compare the political leaning of domains shared by users and their estimated ideology (see Fig. S1 of SM). Results highlight that support for the supply of weapons to Ukraine tends to drop for more extreme political leanings. However, there is still an important distinction between extreme left and extreme right, as the former does have a wider distribution, but the majority of its users are still in favor of military aid, while the latter opposes such aid in its majority.

Coordinated activities

We conclude our analysis by examining how coordinated accounts relate to Active Engagement (AE). Since our data includes only aggregate statistics on view counts, it does not allow us to directly assess the impact of coordinated or bot-like behavior on AE. However, we used retweet information to identify signs of coordination and studied how AE varies in relation to coordination. We classified accounts into two categories, 'Coordinated' and 'Non-Coordinated,' based on whether they retweeted the same content at least 10 times within a 10-s window. Of the 3,181,029 accounts in our dataset, 3.66% were involved in coordinated activities. Moreover, among influencer accounts, 6 out of 49 who opposed aid to Ukraine participated in coordinated activities, while 134 out of 355 influencers who supported Ukraine aid were involved. Additionally, 32.57% of the total 12,475,527 retweets originated from accounts engaged in these coordinated efforts.

Figure 5 presents the distributions of AE for content generated by coordinated (red) and non-coordinated (green) accounts for each action type: retweet(I), reply(II), like(III), and quotes(IV).

Content generated by coordinated accounts shows a higher peak and a narrower distribution compared to non-coordinated content for retweet, reply, and quote actions. Non-coordinated accounts exhibit a broader range of engagement, with the retweet and quote distributions slightly skewed toward higher AE values. The like distributions are more similar, though coordinated accounts display a more pronounced peak. Notably, coordinated accounts exhibit a small secondary peak at high AE values for retweets, which may suggest the presence of a highly coordinated group retweeting the same content.

Overall, the differences in AE between content generated by coordinated and non-coordinated accounts are minor. Results show that while non-coordinated accounts tend to have a wider AE distribution and are slightly more likely to attract active engagement, the similarity in the consumption dynamics between the two groups remains high, suggesting that content from coordinated accounts does not have an increased AE.

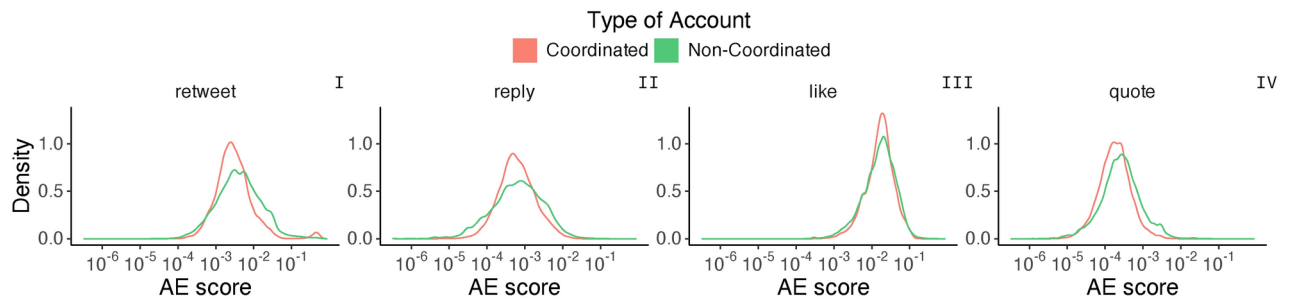


Fig. 5. Distribution of Active Engagement (AE) for content created by coordinated (red) and non-coordinated (green) accounts across each action. While non-coordinated accounts tend to display a broader AE range, coordinated accounts show a small secondary peak at high AE values for retweets.

Discussion

Our study introduces a novel approach, which relies on the recently introduced “impressions count” metric that quantifies content views count, to examine the role of active consumption in online debates and emphasize the factors that can encourage active participation. Complementary to recent research that investigates content visibility and user engagement by considering only active actions, our work, which also incorporates information on passive consumption, focuses on estimating the prominence of active consumption, defined as the proportion of users who actively engage with a tweet through explicit feedback actions after being exposed to content. Given the presence of algorithmic curation that often prioritizes highly engaging content, active consumption plays a significant role in shaping online debates. However, its prevalence is often challenging to quantify due to the lack of proxies for passive consumption, leading it to be overlooked in social media studies.

By comparing the number of active actions with the number of views for tweets related to the Russo-Ukrainian conflict, we analyze the dependency of active participation on several factors at both the user and content levels. Our findings reveal that active actions constitute a smaller portion compared to the number of views, indicating that passive consumers account for a significant share of users who view the content. Moreover, actions that require writing efforts can experience a lower level of active participation than less demanding actions, such as liking. Notably, the most characteristic feature associated with an increased active consumption is the presence of far-right and misinformation-spreading news sources. These contents exhibited a higher ratio of actions per view (i.e. Active Engagement) compared to content from less biased and more reliable sources. Our results complement the existing literature suggesting that polarizing and highly biased political messages are associated with increased engagement⁶⁷ by showing that such content also increases the proportion of users who engage in active feedback actions. In contrast, the author’s popularity, the presence of a coordinated community of accounts, as well as the ideological leaning of the content exhibited a milder effect on active participation. Our results show that, although popular users reach a broader audience, they fail to elicit more active actions from users who see their content, as active engagement has a slight tendency to decrease with users’ popularity. Secondly, since the fraction of active consumers does not depend on the producers’ opinions, the level of active consumption is independent of the echo chamber effect. This suggests that both active and passive consumption dynamics are similar among communities and the level of active engagement inside different echo chambers is similar. Finally, we can speculate on the reasons for the higher active participation obtained by less reliable news sources. The increased active engagement generated by news outlets spreading misinformation may be a consequence of their publishing style. Unreliable sources may be more likely to use emotionally loaded words and pursue clickbait, which may, in turn, trigger more active actions.

Although this study provides an initial understanding of the prevalence of passive consumption in online debate, it also raises several unresolved questions that warrant further investigations. Exploring the prevalence of active consumption in other debates and social networks, such as Facebook or Reddit, is crucial to clarifying its interplay with different environments. Indeed, Twitter is often heavily polarized around topics related to elections⁶⁹, healthcare⁷⁰, or climate change³. Other platforms with lower levels of polarization, such as Reddit⁷¹, may exhibit a different prevalence of active consumption. Understanding the relationship between active consumption and polarization at the platform level is fundamental for assessing its impact on shaping online debates. Further, the study of active consumption in other online environments can shed light on its interplay with several factors influencing online debates such as coordinated behavior and suggestion algorithms. Moreover, Twitter does not provide data about deleted content. Although the amount of deleted content is unlikely to heavily affect our results^{3,72}, analyzing the level of active consumption of deleted content, which may contain highly politicized and extreme messages^{73,74}, could provide interesting insights for designing strategies to identify potentially problematic content. Finally, views count does not provide any information at the individual user level. Studying the characteristics of users who are exposed to content but do not actively interact with it can clarify the role of the interaction structure on information exposure, reveal how content spreads on the platform, and clarify the presence and influence of automated and coordinated accounts in boosting active engagement.

By recognizing the importance of passive consumption alongside more active forms, and examining the relationship between levels of active engagement and factors such as echo chambers, misinformation sources, political bias, coordinated accounts, and other influences on online discussions, this study offers a more comprehensive understanding of engagement dynamics in online social networks, providing a complementary

perspective that is often overlooked in the literature. Unveiling the prevalence of active participation in online debates is fundamental for assessing how it influences suggestion algorithms and, consequently, information spreading. Additionally, studying the content characteristics that can increase active engagement may aid in designing more effective communication strategies to actively involve a higher proportion of users, thereby enlarging the potential outreach of the message. Finally, our work emphasizes the need for comprehensive access to social network data, including a rich set of features and metrics for effective and thorough analysis of online discussions.

Methods

We utilize two datasets—Twitter data and Media Bias/Fact Check (MBFC) domain labels—for our case study of the debate—whether or not to provide military aid to Ukraine in the 2023 war.

Twitter data

We collected tweets regarding Ukraine military aid through the standard Twitter API. We use a keyword-based search query that picks up tweets mentioning one of the following terms: “military aid”, “military support”, “tanks”, “abrams”, “leopard”, “challenger”, “jet”, “aircraft”, “munitions”, “HIMARS”, “rockets” and “missile”. Focusing on a four-month period from November 22nd, 2022 to March 1st, 2023, we retrieved a dataset of 17 M+ tweets posted by 5.2 M+ users. Each of these tweets includes the impression count, a new Twitter feature rolled out in December 2022, which we use to measure the proportion of hidden audience across various means of interaction (i.e., retweet, reply, quote, like) provided by Twitter.

Media bias fact check dataset

To categorize tweets that share a link to news domains we rely on labels retrieved from Media Bias/Fact Check (MBFC)⁶⁸, an independent fact-checking organization that classifies news outlets based on their reliability and political leaning. The MBFC label list we retrieved contains a total of 2190 news outlets, each of which has a domain URL, a political leaning label among “extreme left”, “left”, “left center”, “center”, “right center”, “right”, “extreme right”, and a reliability label among “very low”, “low”, “mixed”, “mostly factual”, “high”, “very high”. For domains with missing labels of ideological leanings, we manually checked the domain page displayed on MBFC to ensure a comprehensive analysis. Using this dataset, we were able to classify 26.6% of all links. We analyzed the most used domains and verified that the unclassified domains are not associated with political or news sources (see Figs. S6, S7 of SM).

Data filtering

The process of data filtering contains the following three steps. First, we filter out tweets that do not have a valid impression count. Intuitively, the impression count should always be greater or equal to any of the active action metrics (i.e., counts of likes, replies, quotes, or retweets). Second, we restrict our analysis to tweets in English to avoid conflating factors (such as geography) that may affect the detection of users’ ideological stances. Third, we only include original tweets in our estimation of hidden audiences, as we do not have information on how the impression count would be passed along downstream in retweets, replies, or quotes.

Interaction network

Using the dataset described in Section Data Filtering, we build the retweet interaction network. This methodology aligns with prevailing practices in Twitter analysis research^{3,5,75}, as retweets are regarded as endorsements of content. On the other hand, quotes and replies are disregarded since they are less likely to signify endorsement and are often used for expressing opposing viewpoints or engaging in polemics⁵. Considering retweets of English content, we build a network by assigning a node to each unique user in the dataset. The set of nodes of the network includes every user who either authored an original English tweet or retweeted an English tweet containing the specified keywords. We create a directed edge from node A to node B if user A retweeted a post authored by user B, and the weight of the edges is determined by the count of unique retweets between the two users, reflecting the strength of their interaction. The final interaction network counts 2.5 million nodes and 7.1 million edges.

Latent ideology estimation

To estimate the ideological stance of users in the debate, we start from the latent ideology algorithm proposed in Refs.^{8,76}. Following the studies already conducted in this field^{3,5}, we consider retweets instead of follower/following relationships as interaction. The latent ideology algorithm requires the extraction of a subset of the influencer which critically affects the ideology estimation results. The method by which such extraction is performed is the main topic of the following subsection. Once the influencer set is known, we apply the Correspondence Analysis algorithm⁷⁷, which follows three steps: (i) Construction of the interaction matrix A , (ii) normalization of the matrix, and (iii) singular value decomposition. For the first step, we construct a matrix A , whose elements A_{ij} represent the number of retweets user i directs toward influencer j . Once A is known, we normalize it as follows. First, we divide by the total number of retweets, obtaining:

$$P = \frac{A}{\sum_{ij} A_{ij}}. \quad (1)$$

Then, we define the following quantities:

$$\begin{cases} r = P1, \\ c = 1^T P, \\ D_r = \text{diag}(r), \\ D_c = \text{diag}(c), \end{cases} \quad (2)$$

and we perform the following normalization operation:

$$S = D_r^{-1/2}(P - rc)D_c^{-1/2}. \quad (3)$$

For the third step, we perform a singular value decomposition of the form $S = U\Sigma V^T$, where U , V are orthogonal matrices and Σ is a diagonal matrix containing the singular values of S . We estimate the ideological leaning of users by taking the subspace associated with the first singular value of the decomposition. Thus, we take the latent ideology the latent ideology of user i to be the i -th entry of the first column of the orthogonal matrix U , while the influencer ideology is estimated by taking the median of its retweeters' ideology scores.

Influencers selection

To apply the ideology-scoring algorithm we first need to extract a set of influencers from the retweet network. The influencer group encompasses several subgroups: (i) Russian and Ukrainian politicians, (ii) official accounts from information media sources such as journals and TV channels, and (iii) political activists. Users in the retweet network are ranked according to their in-degree, corresponding to the number of unique users who have retweeted them. This enables us to start from a manually selected set of prominent accounts pertaining to the three aforementioned categories with some of the highest in-degree. This set then serves as a seed as we further select similar accounts using the “Who to follow” recommendations made by Twitter on their accounts' page. We repeated this procedure until no new accounts were suggested and then refined the selection by excluding accounts with an in-degree lower than 100 and those whose content is unrelated to the Ukrainian conflict. These criteria yield a comprehensive set of 204 influencers, representing both supporters and opponents of military assistance to Ukraine.

Estimation of the active consumption level

The estimation of the active consumption leverages tweet-level metrics, including the number of likes, replies, retweets, quote tweets, and, crucially, the number of views, quantified by the impressions count. We define the level of active consumption as the ratio between the count of an active action—namely, liking, replying, retweeting, and quote retweeting—out of the number of impressions received by a given tweet, a given user, or a given domain, depending on the comparison unit of interest. We call this ratio *Active Engagement (AE)*.

$$AE = \frac{\text{\#of actions}}{\text{\#of impressions}}. \quad (4)$$

Coordinated behavior

We identify coordinated behavior on social media using CoorTweet⁷⁸ to detect patterns where multiple accounts share the same content within a certain timeframe. An action on social media is defined as an account x posting content c at a specific time t . Two or more accounts are considered coordinated when they perform the same action at least n times within a set time window τ . These actions can include sharing the same URL, using the same hashtag, or retweeting the same post. The shared content, or “object”, is tracked as a set of possible coordinated actions $A = \{o_1, o_2, \dots, o_n\}$. Coordination between two accounts x_1 and x_2 occurs when their posts c_1 and c_2 share the same object o , and the time difference $\Delta t = |t_1 - t_2|$ is less than or equal to the time window τ , i.e., $\Delta t \leq \tau$. In our case, we apply this procedure on retweet data, choosing $n = 10$ and $\tau = 10$ as parameters for coordination detection.

Data availability

The code and Twitter data used in this research, shared in accordance with the platform's terms and conditions, are publicly accessible at <https://osf.io/5m3vr/>. MediaBias/FactCheck data is publicly available on the MBFC website: <https://mediabiasfactcheck.com/>.

Received: 6 November 2024; Accepted: 24 January 2025

Published online: 10 February 2025

References

- Andris, C. et al. The rise of partisanship and super-cooperators in the US house of representatives. *PLoS ONE* **10**, e0123507. <https://doi.org/10.1371/journal.pone.0123507> (2015).
- Neal, Z. P. A sign of the times? Weak and strong polarization in the US Congress, 1973–2016. *Soc. Netw.* **60**, 103–112. <https://doi.org/10.1016/j.socnet.2018.07.007> (2020).
- Falkenberg, M. et al. Growing polarization around climate change on social media. *Nat. Clim. Change* **12**, 1114–1121. <https://doi.org/10.1038/s41558-022-01527-x> (2022).
- Global Risks Report. *World Economic Forum*. <https://www.weforum.org/publications/global-risks-report-2024/digest> (2024).
- Flamino, J. et al. Political polarization of news media and influencers on twitter in the 2016 and 2020 us presidential elections. *Nat. Hum. Behav.* **1**, 1–13 (2023).
- Axelrod, R., Daymude, J. J. & Forrest, S. Preventing extreme polarization of political attitudes. *Proc. Natl. Acad. Sci.* **118**, e2102139118 (2021).
- Bramson, A. et al. Disambiguation of social polarization concepts and measures. *J. Math. Sociol.* **40**, 80–111 (2016).

8. Barberá, P. Birds of the same feather tweet together: Bayesian ideal point estimation using twitter data. *Polit. Anal.* **23**, 76–91. <https://doi.org/10.1093/pan/mpu011> (2015).
9. Peralta, A. F., Ramaciotti, P., Kertész, J. & Iniguez, G. Multidimensional political polarization in online social networks. *Phys. Rev. Res.* **6**, 013170 (2024).
10. Cinelli, M., De Francisci Morales, G., Galeazzi, A., Quattrociocchi, W. & Starnini, M. The echo chamber effect on social media. *Proc. Natl. Acad. Sci.* **118**, e2023301118. <https://doi.org/10.1073/pnas.2023301118> (2021).
11. Yarchi, M., Baden, C. & Kligger-Vilenchik, N. Political polarization on the digital sphere: A cross-platform, over-time analysis of interactional, positional, and affective polarization on social media. *Polit. Commun.* **38**, 98–139 (2021).
12. Adamic, L. A. & Glance, N. The political blogosphere and the 2004 US election: Divided they blog. In *Proceedings of the 3rd International Workshop on Link Discovery* 36–43 (2005).
13. Baqir, A., Galeazzi, A., Drocco, A. & Zollo, F. Social media polarization reflects shifting political alliances in Pakistan. Preprint at <http://arxiv.org/abs/2309.08075> (2023).
14. Cinelli, M., Cresci, S., Galeazzi, A., Quattrociocchi, W. & Tesconi, M. The limited reach of fake news on Twitter during 2019 European elections. *PLoS ONE* **15**, e0234689. <https://doi.org/10.1371/journal.pone.0234689> (2020).
15. Vosoughi, S., Roy, D. & Aral, S. The spread of true and false news online. *Science* **359**, 1146–1151. <https://doi.org/10.1126/science.aap9559> (2018).
16. Zhao, Z. et al. Fake news propagates differently from real news even at early stages of spreading. *EPJ Data Sci.* **9**, 7. <https://doi.org/10.1140/epjds/s13688-020-00224-z> (2020).
17. Juul, J. L. & Ugander, J. Comparing information diffusion mechanisms by matching on cascade size. *Proc. Natl. Acad. Sci.* **118**, e2100786118. <https://doi.org/10.1073/pnas.2100786118> (2021).
18. Lazer, D. M. J. et al. The science of fake news. *Science* **359**, 1094–1096. <https://doi.org/10.1126/science.aao2998> (2018).
19. Pariser, E. *The Filter Bubble: What the Internet is Hiding from You* (Penguin, 2011).
20. Ross Arguedas, A., Robertson, C., Fletcher, R. & Nielsen, R. Echo chambers, filter bubbles, and polarisation: A literature review. *Reuters Inst. Study J.* **1**, 1 (2022).
21. Flaxman, S., Goel, S. & Rao, J. M. Filter bubbles, echo chambers, and online news consumption. *Public Opin. Q.* **80**, 298–320 (2016).
22. Tokita, C. K., Guess, A. M. & Tarnita, C. E. Polarized information ecosystems can reorganize social networks via information cascades. *Proc. Natl. Acad. Sci.* **118**, e2102147118. <https://doi.org/10.1073/pnas.2102147118> (2021).
23. Törnberg, P. Echo chambers and viral misinformation: Modeling fake news as complex contagion. *PLoS ONE* **13**, e0203958. <https://doi.org/10.1371/journal.pone.0203958> (2018).
24. Cinelli, M. et al. Selective exposure shapes the Facebook news diet. *PLoS ONE* **15**, e0229129. <https://doi.org/10.1371/journal.pone.0229129> (2020).
25. Bruns, A. Echo chamber? What echo chamber? Reviewing the evidence. In *6th Biennial Future of Journalism Conference (FOJ17)* (2017).
26. Kwak, N., Lane, D. S., Weeks, B. E., Kim, D. H. & Lee, S. S. Now we're talking? Understanding the interplay between online selective and incidental exposure and their influence on online cross-cutting political discussion. *Soc. Sci. Comput. Rev.* **40**, 579–597 (2022).
27. Bruns, A. *Are Filter Bubbles Real?* (Wiley, 2019).
28. Dahlgren, P. M. A critical review of filter bubbles and a comparison with selective exposure. *Nordicom Rev.* **42**, 15–33 (2021).
29. Terren, L. T. L. & Borge-Bravo, R. B.-B. R. Echo chambers on social media: A systematic review of the literature. *Rev. Commun. Res.* **9**, 1 (2021).
30. Gilbert, E., Bergstrom, T. & Karahalios, K. Blogs are echo chambers: 42nd Annual Hawaii International Conference on System Sciences, HICSS. In *Proceedings of the 42nd Annual Hawaii International Conference on System Sciences, HICSS*. <https://doi.org/10.1109/HICSS.2009.91> (2009).
31. Edwards, A. (How) do participants in online discussion forums create echo chambers?: The inclusion and exclusion of dissenting voices in an online forum about climate change. *J. Argument. Context* **2**, 127–150. <https://doi.org/10.1075/jaic.2.1.06edw> (2013).
32. Del Vicario, M. et al. Echo chambers: Emotional contagion and group polarization on facebook. *Sci. Rep.* **6**, 37825. <https://doi.org/10.1038/srep37825> (2016).
33. Cossard, A. et al. Falling into the echo chamber: The Italian vaccination debate on twitter. *Proc. Int. AAAI Conf. Web Soc. Media* **14**, 130–140 (2020).
34. Santoro, A. et al. Analyzing the changing landscape of the covid-19 vaccine debate on twitter. *Soc. Netw. Anal. Min.* **13**, 115 (2023).
35. Del Vicario, M. et al. The spreading of misinformation online. *Proc. Natl. Acad. Sci.* **113**, 554–559. <https://doi.org/10.1073/pnas.1517441113> (2016).
36. Morini, V., Pollacci, L. & Rossetti, G. Toward a standard approach for echo chamber detection: Reddit case study. *Appl. Sci.* **11**, 5390 (2021).
37. Buongiovanni, C. et al. Will you take the knee? Italian twitter echo chambers' genesis during euro 2020. In *International Conference on Complex Networks and Their Applications* 29–40 (Springer, 2022).
38. Garimella, K., Morales, G. D. F., Gionis, A. & Mathioudakis, M. Quantifying controversy on social media. *ACM Trans. Social Comput.* **1**, 1–27. <https://doi.org/10.1145/3140565> (2018).
39. Diaz-Diaz, F., San Miguel, M. & Meloni, S. Echo chambers and information transmission biases in homophilic and heterophilic networks. *Sci. Rep.* **12**, 9350. <https://doi.org/10.1038/s41598-022-13343-6> (2022).
40. Hohmann, M., Devriendt, K. & Coscia, M. Quantifying ideological polarization on a network using generalized Euclidean distance. *Sci. Adv.* **9**, 2044. <https://doi.org/10.1126/sciadv.abq2044> (2023).
41. Baumann, F., Lorenz-Spreen, P., Sokolov, I. M. & Starnini, M. Modeling echo chambers and polarization dynamics in social networks. *Phys. Rev. Lett.* **124**, 048301. <https://doi.org/10.1103/PhysRevLett.124.048301> (2020).
42. Baumann, F., Lorenz-Spreen, P., Sokolov, I. M. & Starnini, M. Emergence of polarized ideological opinions in multidimensional topic spaces. *Phys. Rev. X* **11**, 011012. <https://doi.org/10.1103/PhysRevX.11.011012> (2021).
43. Klapper, J. T. *The Effects of Mass Communication* (Free Press, 1960).
44. Nickerson, R. S. Confirmation bias: A ubiquitous phenomenon in many guises. *Rev. Gen. Psychol.* **2**, 175–220. <https://doi.org/10.1037/1089-2680.2.2.175> (1998).
45. Brown, M. A. et al. *Echo Chambers, Rabbit Holes, and Algorithmic Bias: How Youtube Recommends Content to Real Users* (2022).
46. Ciampaglia, G. L., Nematzadeh, A., Menczer, F. & Flammini, A. How algorithmic popularity bias hinders or promotes quality. *Sci. Rep.* **8**, 15951 (2018).
47. Bakshy, E., Messing, S. & Adamic, L. A. Exposure to ideologically diverse news and opinion on facebook. *Science* **348**, 1130–1132 (2015).
48. Shao, C. et al. The spread of low-credibility content by social bots. *Nat. Commun.* **9**, 1–9 (2018).
49. Stella, M., Ferrara, E. & De Domenico, M. Bots increase exposure to negative and inflammatory content in online social systems. *Proc. Natl. Acad. Sci.* **115**, 12435–12440 (2018).
50. Cresci, S. A decade of social bot detection. *Commun. ACM* **63**, 72–83 (2020).
51. Tardelli, S. et al. Temporal dynamics of coordinated online behavior: Stability, archetypes, and influence. *Proc. Natl. Acad. Sci.* **121**, e2307038121 (2024).
52. Gong, W., Lim, E.-P. & Zhu, F. Characterizing silent users in social media communities. *Proc. Int. AAAI Conf. Web Soc. Media* **9**, 140–149. <https://doi.org/10.1609/icwsm.v9i1.14582> (2015).

53. Antelmi, A., Malandrino, D. & Scarano, V. Characterizing the øthe 90-9-1 rule. In *Companion Proceedings of The 2019 World Wide Web Conference, WWW '19*, 1035–1038. <https://doi.org/10.1145/3308560.3316705> (Association for Computing Machinery, 2019).
54. Gong, W., Lim, E.-P. & Zhu, F. Characterizing silent users in social media communities. *Proc. Int. AAAI Conf. Web Soc. Media* **9**, 140–149 (2015).
55. Tagarelli, A. & Interdonato, R. “Who’s out there?” Identifying and ranking lurkers in social networks. In *Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining* 215–222 (2013).
56. Perna, D., Interdonato, R. & Tagarelli, A. Identifying users with alternate behaviors of lurking and active participation in multilayer social networks. *IEEE Trans. Comput. Soc. Syst.* **5**, 46–63 (2017).
57. Caliò, A., Interdonato, R., Pulice, C. & Tagarelli, A. Topology-driven diversity for targeted influence maximization with application to user engagement in social networks. *IEEE Trans. Knowl. Data Eng.* **30**, 2421–2434 (2018).
58. Power, P., Kubey, R. & Kiousis, S. Audience activity and passivity: An historical taxonomy. *Ann. Int. Commun. Assoc.* **26**, 116–159 (2002).
59. Levy, M. R. & Windahl, S. Audience activity and gratifications: A conceptual clarification and exploration. *Commun. Res.* **11**, 51–78 (1984).
60. Page, B. I., Shapiro, R. Y. & Dempsey, G. R. What moves public opinion? *Am. Polit. Sci. Rev.* **81**, 23–43 (1987).
61. Boulianne, S. & Larsson, A. O. Engagement with candidate posts on twitter, instagram, and facebook during the 2019 election. *New Media Soc.* **25**, 119–140 (2023).
62. Galeazzi, A., Peruzzi, A., Brugnoli, E., Delmastro, M. & Zollo, F. Unveiling the hidden agenda: Biases in news reporting and consumption. *PNAS Nexus* **3**, 474 (2024).
63. Papakyriakopoulos, O., Serrano, J. C. M. & Hegelich, S. Political communication on social media: A tale of hyperactive users and bias in recommender systems. *Online Soc. Netw. Media* **15**, 100058 (2020).
64. Vilella, S., Semeraro, A., Paolotti, D. & Ruffo, G. Measuring user engagement with low credibility media sources in a controversial online debate. *EPJ Data Sci.* **11**, 29 (2022).
65. Li, Y. & Xie, Y. Is a picture worth a thousand words? An empirical study of image content and social media engagement. *J. Mark. Res.* **57**, 1–19 (2020).
66. *Twitter API Documentation*. <https://developer.twitter.com/en/docs/twitter-ads-api/creatives/api-reference/media-library>.
67. Tucker, J. A. et al. Social media, political polarization, and political disinformation: A review of the scientific literature. In *Political Polarization, and Political Disinformation: A Review of the Scientific Literature (March 19, 2018)* (2018).
68. Media Bias/Fact Check. <https://mediabiasfactcheck.com/> (2023).
69. Garimella, V. R. K. & Weber, I. A long-term analysis of polarization on twitter. *Proc. Int. AAAI Conf. Web Soc. Media* **11**, 528–531 (2017).
70. Jiang, X. et al. Polarization over vaccination: Ideological differences in twitter expression about covid-19 vaccine favorability and specific hesitancy concerns. *Soc. Media Soc.* **7**, 205630512110484 (2021).
71. De Francis Morales, G., Monti, C. & Starnini, M. No echo in the chambers of political interactions on reddit. *Sci. Rep.* **11**, 2818 (2021).
72. Di Bona, G. et al. Sampled datasets risk substantial bias in the identification of political polarization on social media. Preprint at <http://arxiv.org/abs/2406.19867> (2024).
73. Bastos, M. This account doesn’t exist: Tweet decay and the politics of deletion in the brexit debate. *Am. Behav. Sci.* **65**, 757–773 (2021).
74. Rogers, R. Deplatforming: Following extreme internet celebrities to telegram and alternative social media. *Eur. J. Commun.* **35**, 213–229 (2020).
75. Conover, M. et al. Political polarization on twitter. *Proc. Int. AAAI Conf. Web Soc. Media* **5**, 89–96 (2011).
76. 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. <https://doi.org/10.1177/0956797615594620> (2015).
77. Greenacre, M. J. Correspondence analysis. *Wiley Interdiscip. Rev. Comput. Stat.* **2**, 613–619 (2010).
78. Righetti, N. & Balluff, P. *Coortweet: Coordinated Networks Detection on Social Media* (2024).

Acknowledgements

This work is the revised output of the Complexity72h workshop, held at the IFIS in Palma, Spain, 26–30 June 2023, see <https://www.complexity72h.com>. A.G. thanks CY4GATE and the National Recovery and Resilience Plan (NRRP) Project “Securing software Platforms - SOP”, CUP H73C22000890001 for the financial support. A.G. is grateful to Max Falkenberg, Fabiana Zollo and Matteo Cinelli for the fruitful discussions and comments.

Author contributions

A.G. conceived the research and retrieved the data. A.G. and M.T. supervised the analysis. A.B., Y.C., F.D.D., S.K., T.L., V.M., and V.P. analyzed the data and visualized the results. All authors participated in the discussion of the results, contributed to the writing, and revised 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/10.1038/s41598-025-88117-x>.

Correspondence and requests for materials should be addressed to A.G.

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