

<https://doi.org/10.1038/s44260-024-00008-w>

Affective polarization and dynamics of information spread in online networks



Kristina Lerman^{1,2}✉, Dan Feldman², Zihao He² & Ashwin Rao²

Members of different political groups not only disagree about issues but also dislike and distrust each other. While social media can amplify this emotional divide—called affective polarization by political scientists—there is a lack of agreement on its strength and prevalence. We measure affective polarization on social media by quantifying the emotions and toxicity of reply interactions. We demonstrate that, as predicted by affective polarization, interactions between users with same ideology (in-group replies) tend to be positive, while interactions between opposite-ideology users (out-group replies) are characterized by negativity and toxicity. Second, we show that affective polarization generalizes beyond the in-group/out-group dichotomy and can be considered a structural property of social networks. Specifically, we show that emotions vary with network distance between users, with closer interactions eliciting positive emotions and more distant interactions leading to anger, disgust, and toxicity. Finally, we show that similar information exhibits different dynamics when spreading in emotionally polarized groups. These findings are consistent across diverse datasets spanning discussions on topics such as the COVID-19 pandemic and abortion in the US. Our research provides insights into the complex social dynamics of affective polarization in the digital age and its implications for political discourse.

Democrats and Republicans in the United States not only disagree on many economic, political and cultural issues but have also grown less tolerant of opposing viewpoints, with members of each party disliking and distrusting those affiliated with the opposing party. The emotional divide—dubbed *affective polarization* by political scientists^{1,2}—has become a destabilizing force in a democracy, reducing cooperation across party lines, stoking hostility towards out-group members^{3–6} and eroding trust in experts and institutions. For example, the polarized response to the COVID-19 pandemic led individuals affiliated with one political party to distrust recommendations from public health experts if they were supported by the opposing party, hindering effective response to the health crisis^{7–9}.

Research shows that affective polarization is driven by factors including the news media, political elites, and demographics^{2,4,10}. Although scholars disagree on how much social media contributes to partisan animosity¹¹, many see it as an important amplifier¹². Social media discourse tends to promote inflammatory language and moral outrage directed at the out-group^{13–15}. Moreover, social media echo chambers, which segregate users within communities of like-minded others, may amplify political polarization by exposing users to extreme and divisive content¹⁶. Beyond echo-chambers, exposure to out-party views exists¹⁷ and may worsen political polarization¹⁸.

This paper describes an empirical study that makes a three-fold contribution to this body of research: we propose an instrument to quantify

affective polarization in social media, establish an empirical relationship between emotions and structure of online networks, and characterize partisan differences in the dynamics of information spread on them. In contrast to existing research, which measures affective polarization by looking at how people talk *about* out-group members^{15,19–21}, we focus on how they talk *to* them. We leverage state-of-the-art language models to measure emotions and toxicity of online interactions between social media users and demonstrate they exhibit the hallmarks of affective polarization¹, namely *in-group favoritism–out-group animosity*. Moreover, these emotions vary with distance between interacting users in a social network, demonstrating an association between affective polarization and the structure of online networks. This idea is illustrated in Fig. 1: when replying to people closer to them in a network (e.g., a retweet network), users express more positive emotions, but when replying to those who are farther away, they express more negative emotions and toxicity. These findings are consistent across datasets of discussions about the COVID-19 pandemic²² and abortion²³ on Twitter. Finally, we analyze the spread of information and show that discussions of contentious issues within partisan groups exhibit different dynamics. Some issues show random bursts of re-sharing, consistent with being driven by the news cycle, while others persist over longer periods of time, reflecting how their emotional salience to ideological divisions helps focus attention of polarized groups.

¹USC Information Sciences Institute, Marina del Rey, CA, USA. ²USC, Computer Science Dept, Los Angeles, CA, USA. ✉e-mail: lerman@isi.edu




Fig. 1 | Affective polarization in networks. People express warmer feelings when replying to those closer to them in a social network; when replying to those farther away, they express more negativity.

Table 1 | Number of retweets and replies from the datasets in our study

Dataset	#Retweets	#Retweeters	#Replies	#Responders
Roe_v_Wade	7,131,980	1,005,156	460,868	172,988
COVID-19	46,419,871	10,758,690	4,173,679	833,875

Number of retweeters and the number of responders gives the total number of unique user ids that participated in a retweet and reply interaction respectively.

Our study sheds light on the complex dynamics of affective polarization in the digital age, offering insights into the emotional foundations of political discourse on social media and the interaction of emotions with network structure and information diffusion.

Results

We study two massive datasets of online discussions on Twitter: tweets about the *COVID-19* pandemic in the US and tweets about the overturning of the *Roe v Wade* decision that legalized abortion in US in 1973. We classify users' ideology as liberal or conservative based on the text of their tweets, and use transformer-based language models to detect emotions and toxic language in their replies (see Methods). Finally, we identify contentious issues in online discussions and study dynamics of re-sharing by each ideological group. Table 1 summarizes statistics of the two datasets.

Emotions of In-group vs Out-group Interactions

We measure affective polarization by quantifying emotions expressed in reply interactions between users with a known ideology. *In-group interactions* are replies between same-ideology users: liberal replying to a liberal, etc. *Out-group interactions* are replies between opposite-ideology users: liberal replying to a conservative or vice versa. Figure 2 shows the distribution of emotion confidence scores of replies in the abortion dataset. Anger, disgust and toxicity scores of out-group replies are substantially higher than for in-group replies, consistent with *out-group animosity*. Similarly, in-group replies have higher scores for joy, consistent with *in-group favoritism*. We observe similar trends in the COVID-19 data (Supplementary Fig. 1). All differences are statistically significant. The differences between in-group and out-group interactions remain after disaggregating replies by ideology across both datasets (Supplementary Fig. 2 and Supplementary Fig. 3).

Our empirical analysis yields additional insights. We did not find significant differences in love, sadness or optimism. Out-group interactions are more emotional but shorter. Additionally, fear does not behave like other negative emotions in that it is higher for in-group interactions. This may reflect the role of fear as an agent of social cohesion. By identifying a

common threat (see folktales and legends^{24,25}), fear creates solidarity, which unites the group.

Figure 3 summarizes these results for both datasets by showing the difference between the mean emotion confidence of out-group and in-group replies. There exists a consistent emotion gap: interactions across ideological lines exhibit more negativity and toxicity, and less joy. Surprisingly, out-group replies are significantly shorter: in the Roe_v_Wade data, out-group replies are on average 17 characters shorter, while in the COVID-19 data, they are 8 characters shorter. Together with the finding that out-group interactions are more emotional, this suggests that people communicate across party lines not to convey information but rather to express animosity and to troll.

Emotions and network distance

Next, we explore the interplay between emotions and network structure, using retweet networks to represent the online social networks (see Methods). Due to their large size, we use an embedding technique *LargeVis*²⁶ to visualize the networks. Figure 4 shows the heatmap of these embeddings, with bright spots corresponding to dense clusters of users who frequently retweet one another.

The retweet network of the discussions in the Roe_v_Wade data (Fig. 4a) shows two main clusters with additional substructure near the larger cluster. The coarse-grained structure reflects the polarized nature of the abortion debate: most of the liberals are in the larger cluster and the conservatives are in the smaller one. The retweet network of the COVID-19 discussions (Fig. 4b) contains many small clusters. Although these discussions did grow to be polarized, during the first months of the pandemic covered by our dataset, these divisions do not appear entrenched.

We measure network distance between two users in the retweet network embedding space (see Methods). Figure 5 shows confidence scores of emotions in reply interactions as a function of network distance between interacting users in the Roe_v_Wade data. We divide the distances into four equal-statistics bins, or quartiles. The first quartile (Q1) represents the 25% of the closest pairs of interacting users, while Q4 represents the 25% of the most distant pairs. There are systematic differences in emotions: when users reply to those farther away in the retweet network, they tend to express more anger and disgust, and use more toxic language; while they tend to express more joy and fear in replies to closer users and also use more words. All differences between the quartiles are statistically significant.

To measure these trends, we perform linear regression on emotion confidence scores using network embedding distance as the independent variable (see Supplementary Figs. 4 and 5). Figure 6 brings together all regression coefficients across both datasets, showing similar trends across datasets, in agreement with Fig. 3. Refer to Supplementary Table 1 and Supplementary Table 1 for detailed regressions results. The relationship between emotional expression and network distance holds separately for in-group and out-group interactions and after disaggregating by ideology (Supplementary Figs. 7, 8). These results show that we can measure affective polarization even in the absence of partisan labels that define the out-group. Emotions and networks interact so that users feel warmer towards those who are closer to them and colder towards those who are farther away, regardless of group affiliation.

For robustness, we calculate the shortest path length between pairs of nodes in the retweet network as an alternate measure of network distance. We find that the trends do not depend on how we measure network distance, with results largely remaining the same when using either metric (Supplementary Fig. 8), even after disaggregating by ideology (Supplementary Figs. 9 and 10).

Information spread in affectively polarized populations

Users interact across diverse network distances, seeing information shared by allies and foes alike. Their emotional reactions modulate attention to contentious issues, affecting how those issues spread through the network.

Roughly half of the tweets in the COVID-19 corpus discuss at least one contentious *issue*, such as vaccines, education, masking, lockdowns, or




Fig. 2 | In-group and out-group affect in the Roe_v_Wade data. Boxplots show confidence scores of emotions expressed in replies interactions in the abortion discussions between same-ideology users (in-group) and opposite-ideology users (out-group). Out-group interactions show more **a** anger, **b** disgust, and use more **c** toxic language, but also slightly less **d** joy and **e** fear, and are generally **f** shorter in length (as measured by the number of characters divided by the maximum allowed

length, 280 characters). The boxes span the first to third quartiles, with whiskers extending 1.5 times the interquartile range. The horizontal line inside the box represents the median, and diamond symbol marks the mean. Differences in means were tested for statistical significance using a two-sided Mann-Whitney U Test with the Bonferroni correction: * indicates significance at $p < 0.05$, ** - $p < 0.01$, *** - $p < 0.001$, **** - $p < 0.0001$ and, ns not-significant.




Fig. 3 | Affective polarization. Difference between the mean emotion confidence scores of out-group and in-group interactions in the Roe_v_Wade and COVID-19 datasets. Error bars show standard errors.

origins of the virus (see Methods). Attention to issues waxes and wanes as events drive partisan interest and the diffusion of tweets. Figure 7 (top row) plots the daily number of retweets of each issue by liberals and conservatives at different periods of the pandemic. The absolute number of retweets varies greatly due to difference in the size of the partisan groups. To address the imbalance we standardize the number of retweets within each group using z-score normalization.

Each time series in Fig. 7 represents the complex dynamics of information diffusion within a group. To characterize these dynamics, we

calculate the autocorrelation function (ACF), which measures the correlation of different points in the same time series, separated by time lags. The middle row of Fig. 7 shows ACF along with confidence intervals (blue lines). Some ACF plots show weekly patterns in the volume of retweets with peaks at 7, 14, etc. days. However, others show a rapid or gradual decay of ACF to non-significant values. The former trend is consistent with short-lived spikes in retweets occurring at random times, while the latter trend is indicative of persistent attention. The bottom row of Fig. 7 shows the time lags at which the ACF drops below the confidence interval. The early pandemic (post-President Trump's declaration of national emergency, left column of Fig. 7) was characterized by school closures and challenges of online learning, a topic favored by liberals. Ending lockdowns and reopening the economy became an important issue for conservatives, culminating in protests at state capitals in April 2020. Consistently, discussions about education were persistent among liberals, while discussions about lockdowns were more persistent among conservatives. The summer of 2020 (right column of Fig. 7) saw mass protests for racial justice, sparked by the murder of George Floyd, in which many liberals participated. Since demonstrations required violating social distancing measures enacted to limit the spread of the virus (lockdowns issue), liberals promoted masking to stay safe, which made masking an important issue among both liberals and conservatives.

In June 2022, the Supreme Court reversed federal guarantees on abortion access made by the 1973 Roe v Wade decision. The decision, as well as the leak of its draft in May 2022, sparked debates on all abortion issues: women's health, bodily autonomy, exceptions to abortion bans, religion, and fetal rights. The difference in the nature of information spread is evident

Fig. 4 | Heatmap of the embeddings of social networks. Each network was constructed by linking users who retweeted each other in online discussions about a) the 2022 overturning of Roe_v_Wade and b) the COVID-19 pandemic. The massive retweet networks were embedded in a lower-dimensional space using a graph embedding method. The heatmap of the embedding shows bright spots of densely-linked communities of users who frequently retweet one another. The retweet network of abortion discussions **a** shows two overarching polarized communities, while the network of the pandemic discussions **b** has a multi-focal structure.





Fig. 5 | Affective polarization in the retweet network of Roe_v_Wade data. Each plot shows emotions in reply interactions as a function of network distance between interacting users. Network distance is calculated in the network embedding space and then divided into quartiles, with Q1 showing the 25% of closest interactions and Q4 showing the 25% of the most distant interactions. Emotions **a** anger, **b** disgust and **c** toxicity increase with distance between users in the network embedding space, while **d** joy and **e** fear decrease with distance, as does **f** reply length. The boxes span

the first to third quartiles, with whiskers extending 1.5 times the interquartile range. The horizontal line inside the box represents the median, and diamond symbol marks the mean. Differences in means were tested for statistical significance using a two-sided Mann-Whitney U Test with the Bonferroni correction: * indicates significance at $p < 0.05$, ** - $p < 0.01$, *** - $p < 0.001$, **** - $p < 0.0001$ and, ns - not-significant.

in Fig. 8, which compares the autocorrelation function of the time series of the daily volume of retweets in the 80-day period before the leak and after the decision. Before the leak (left column of Fig. 8), the irregular bursts of retweets, triggered by events, which brought short-lived spikes of attention to issues. After the overturning of Roe v Wade, dynamics of information spread among conservatives changed, with issues related to religion, fetal rights and exceptions to abortion ban reverberating among this group.

Discussion

We investigated the emotional dimension of political polarization in online networks and the interplay between emotions, network structure, partisanship and dynamics of information spread. Analyzing emotions expressed in online discussions about abortion and the COVID-19 pandemic, we found that users expressed more emotions in their replies to opposite-ideology users and their valence had the hallmarks of affective polarization,




Fig. 6 | Regression coefficients of affect as a function of network distance in the Roe_v_Wade and COVID-19 datasets. The bars represent the value of the regression coefficient of the emotion or toxicity of replies as a function of normalized distance between interacting users in the network embedding space. Distance was normalized by rescaling distances in the embedding space to unit interval. Error bars show 95% confidence interval.

namely “in-group favoritism, out-group animosity”¹. Importantly, we showed that affective polarization generalizes beyond the in-group/out-group dichotomy. When accounting for network distance between interacting users in the retweet network, a proxy of the follower graph, anger, disgust and toxicity increased with distance, while joy largely decreased. These findings generalized across datasets and measures of network distance confirming robustness of findings.

Our findings shed more light on affective polarization. Only a subset of emotion and toxicity categories we measured displayed group differences: anger, disgust, fear, joy, toxic language and obscenities. Our study also revealed that like joy, fear is usually higher within in-group replies. This stands in contrast to previous studies²⁷, but highlights the complex role of fear in social interactions. Studies of folklore and mythology suggest that fear helps social cohesion: by making threats salient, fear increases in-group solidarity^{24,26}. In-group replies also tend to be longer, supporting the notion that they serve to share information within the group on how to negotiate threats.

One implication of this finding is that information may spread differently among interacting groups within the same population based on its emotional salience to each group. We saw some evidence for this in how liberal and conservative users shared various issues during the COVID-19 pandemic. Conservatives paid more attention to lockdowns during the early phase of the COVID-19 pandemic, as demonstrated by the persistence of retweets about lockdowns. During this period of time conservatives protested lockdowns, suggesting the emotional importance of this issue in differentiating them from liberals. At other times the lockdowns issue attracted short bursts of attention, triggered by events in the news. Similar patterns in our data point to the complex interplay between emotions, partisanship and dynamics of information spread within a polarized population.

Our results also highlight important differences between reply and retweet interactions. While researchers sometimes conflate them when building social networks, our findings suggest that these interactions serve a very different purpose and that combining them may obfuscate important features of network structure.

Like any study of social media, ours has limitations that temper conclusions. By necessity, our datasets represented a small sample of online discussions, even when controlling for the topic. Replication of results in more representative samples of online discussions could help verify the generalizability of findings. Beyond biases introduced by keyword-centered tweet collection, not observing all interactions may have limited the range of emotions we observed. Similarly, retweets are a biased sample of the follower relationships²⁸, which may have impacted our conclusions. Errors in partisanship detection, emotion and toxicity classification, may have further

affected our findings. While we cannot discount all of these biases, the consistency of our results across datasets and scenarios gives us confidence about our conclusions.

Our study does not disentangle the directionality of the relationship between network distance and emotive expression, limiting its theoretical contributions. Despite this limitation, we believe that the descriptive analysis of the interplay between emotions and network structure is still valuable.

Another thing to consider is that our results could be explained by some confounder, rather than group polarization or network structure. For example, emotionally charged content is retweeted more frequently¹³, adding emotional texture to retweet networks. Moreover, Twitter’s personalization algorithm may highlight emotionally charged content, thereby driving engagement¹⁴, while rapid information spread within communities²⁹ may also distort emotions. Although we cannot discount all confounders, the consistency of our findings across different datasets is encouraging.

Despite potential limitations like incomplete observability and data bias, the consistency of results across datasets and methods provides confidence in the conclusions of our study about the interplay between emotions, network structure and the dynamics of information spread across and within groups. Understanding these mechanisms is crucial for addressing challenges related to misinformation, polarization, and the health of public discourse in the digital age.

Methods

Data

We used a public corpus of tweets about the COVID-19 pandemic²², focusing on tweets posted between January 21, 2020 and April 22, 2020 in the analysis of polarization. Our second dataset is a public corpus of tweets about abortion rights collected between January 1, 2022 to January 6, 2023²³. The tweets contain keywords and hashtags that reflect both sides of the abortion rights debate in US during the period that Roe v Wade was overturned.

For both datasets, Carmen³⁰, a geo-location tool for Twitter data, was used to link tweets to locations within US. Carmen relies on metadata in tweets, such as “place” and “coordinates” objects that encode location, as well as mentions of locations in a user’s bio, to infer their location. We used this approach to filter out users whose home location is not one of US states.

To study dynamics of emotional polarization and information spread, we focus on interpersonal interactions in online social networks. On platforms like Twitter, these engagements predominantly manifest through retweets and replies. Each retweet or reply post typically refers back to an original tweet, hereinafter referred to as the parent tweet. Furthermore, every retweet or reply record includes the ID of its author as well as the ID of the author of the parent tweet. We analyze all retweets and replies where both the author and the referenced tweet’s author are situated in the United States and discard the rest. Table 1 shows statistics of the resulting datasets.

Emotions & toxicity detection

Emotions represent feelings, which are often expressed through language. Early attempts to automate emotion recognition from text relied on emotion lexicons—curated collections of words categorized by their emotional content, e.g., LIWC³¹, EmoLex³², and WKB³³. The advent of transformers has revolutionized emotion detection, which could now benefit from contextual cues.

To measure emotions we use an open-source library Demux³⁴. This model was shown to outperform competing methods on the SemEval 2018 Task 1 e-c benchmark³⁵. Demux assigns none, one or more emotions to input text, along with a scalar value representing its confidence score. The confidence score is the likelihood the tweet expresses that emotion. Demux can recognize a range of emotions in multi-lingual text, including anger, disgust, fear, sadness, joy, love, trust, pessimism and optimism.

To measure toxicity, we use an open-source classifier Detoxify (<https://github.com/unitaryai/detoxify>). The model is trained on the multilabel toxic comment classification task to recognize toxicity levels of tweets. The model outputs a score, a scalar value that captures the likelihood the tweet expresses




Fig. 7 | Persistent dynamics of information spread in COVID-19 discussions. The figures (a, b) show the time series of the volume of retweets on each issue made by liberals and conservatives, (c, d) show the autocorrelation function of each time series, and (e, f) show the longest significant time lag, in days, of the autocorrelation function. Each column represents a different 80 day time period. The column on left

highlights dynamics post declaration of National Emergency and the column on right highlights dynamics after the July 4th holiday weekend. The rows represent pandemic-related issues: education and online learning, vaccines, masking, lockdowns and social distancing, and origins of the coronavirus.




Fig. 8 | Dynamics of information spread in abortion discussions. The figures (a, b) show the longest significant time lag, in days, of the autocorrelation function. Each column represents a different 80 day time period, (a) before the leak of the Dobbs

decision, (b) after after SCOTUS's Dobbs decision. The rows represent abortion issues, such as women's health, bodily autonomy, exceptions to abortion bans, religion, and fetal rights.

toxicity, severe toxicity, obscenity, a threat, or an insult. In this study, we only use toxicity scores.

Ideology classification

To estimate the ideology of social media users, studies have relied on follower relationships³⁶, mention and retweet interactions^{37,38}, and URL sharing^{16,39–41}. Here we use a method described in⁴¹ to classify individual Twitter users as *liberal* or *conservative* based on the text of the messages they share. The classifier leverages political bias scores assigned to well over 6K online sources by Media Bias-Fact Check (MBFC)⁴². Based on these scores, training data is created by assigning each user a score that is a weighted average of the political bias scores of the URLs they share. After training a text embedding-based model on this data, the classifier achieves state-of-the-art performance on recognizing user ideology.

Issue detection

Contentious issues that emerged during the pandemic include (1) the *origins of the virus*, involving debates over bats, wet-markets, lab leak and the gain of function research; (2) the implementation of *lockdown* measures via quarantines, business closures, social distancing and bans on mass congregation; (3) *masking* mandates and face mask shortages; (4) the impact of the pandemic on *education* with school closures and shift to online learning; and (5) *vaccine*-related discussions^{43–47}.

The issues central to the abortion debate in the US^{48,49} include (1) *religion* and faith-based arguments against abortion; (2) views promoting primacy of *fetal rights*; (3) framing abortion as a *bodily autonomy* issue and freedom to choose; (4) abortion as a women's *health* issue; and (5) the question of *exceptions* to abortion restrictions, for example, to save a woman's life or in the case of rape or incest.

We leverage a method developed and validated in previous works to detect these issues. The method harvests relevant keywords and phrases from Wikipedia pages discussing specific issues, labels a subset of tweets using these terms and trains a transformer-based model on this data. The trained models were shown to achieve state-of-the-art performance recognizing pandemic and abortion-related issues in these datasets^{50,51}. A tweet could discuss multiple issues or no issue at all.

Network construction

Studies of online social networks differ in how they represent edges between users. Some researchers^{16,52} use follower relations to capture the attention users pay to others. However, collecting follower links is highly impractical due to API limitations. Instead, researchers rely on retweets, which can be

more easily extracted from the tweets metadata, to construct the social graph²⁸. Retweet networks are foundational to social media analysis and have been used in studies of information spread³³, virality prediction²⁹, fake news⁵⁴, online echo chambers^{16,55}, political polarization^{37,56}, and online discussions^{57,58}. Following this practice, we construct a retweet-based social network for each dataset. Retweets are evidence that both the author of the original tweet and the author of the retweet were, at least on one occasion, exposed to the same information. Therefore, we model retweet networks as undirected, unweighted graphs whose nodes represent users and edges represent the existence of at least one retweet between them (in either direction).

We measure distance in networks in two ways. The first one uses the shortest path between two nodes in the retweet network. The second one measures Euclidean distance between nodes in the embedding space (Fig. 4) generated by the LargeVis model.

Reporting summary

Further information on research design is available in the Nature Research Reporting Summary linked to this article.

Data availability

Given the restrictions imposed by X (then Twitter) on publicly sharing tweet objects, the authors state that there are certain restrictions on its availability. Only the tweet identifiers used in this study are publicly available. Readers can apply for API access from X in order to rehydrate these tweets. The tweet identifiers for the COVID-19 dataset are available at <https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/DKOVLA> and ones for the Roe-Wade dataset are available at <https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/STU0J5>. The DOIs for the COVID-19 and Roe-Wade datasets are: 10.7910/DVN/DKOVLA and 10.7910/DVN/STU0J5 respectively.

Code availability

The ideology classifier used to determine individual ideological leanings in this work is available at <https://tinyurl.com/yy7xxsey>. The emotion classifier is available at: <https://github.com/hasanhuz/SpanEmo>. Toxicity is assessed using the Detoxify model at <https://github.com/unitaryai/detoxify>. Geolocation inference is performed using Carmen described in <https://github.com/mdredze/carmen>. Code for network visualizations were obtained from <https://github.com/lferry007/LargeVis>. Code and auxiliary datasets used as a part of this study are also made available at <https://zenodo.org/records/10810851>.

Received: 11 January 2024; Accepted: 6 May 2024;

Published online: 07 June 2024

References

1. Iyengar, S., Sood, G. & Lelkes, Y. Affect, not ideology: a social identity perspective on polarization. *Public Opin. Q.* **76**, 405–431 (2012).
2. Iyengar, S., Lelkes, Y., Levendusky, M., Malhotra, N. & Westwood, S. J. The origins and consequences of affective polarization in the united states. *Annu. Rev. Political Sci.* **22**, 129–146 (2019).
3. Kingzette, J. et al. How affective polarization undermines support for democratic norms. *Public Opin. Q.* **85**, 663–677 (2021).
4. Whitt, S. et al. Tribalism in america: behavioral experiments on affective polarization in the trump era. *J. Exp. Political Sci.* **8**, 247–259 (2021).
5. Rudolph, T. J. & Hetherington, M. J. Affective polarization in political and nonpolitical settings. *Int. J. Public Opin. Res.* **33**, 591–606 (2021).
6. Bertzen, L. E., Kelsall, H. & Hartevel, E. Consequences of affective polarization: Avoidance, intolerance and support for violence in the united kingdom and norway. *European Journal of Political Research* (2023).
7. Green, J., Edgerton, J., Naftel, D., Shoub, K. & Cranmer, S. J. Elusive consensus: Polarization in elite communication on the covid-19 pandemic. *Sci. Adv.* **6**, eabc2717 (2020).
8. Grossman, G., Kim, S., Rexer, J. M. & Thirumurthy, H. Political partisanship influences behavioral responses to governors' recommendations for covid-19 prevention in the united states. *Proc. Natl Acad. Sci. USA* **117**, 24144–24153 (2020).
9. Gollwitzer, A. et al. Partisan differences in physical distancing are linked to health outcomes during the covid-19 pandemic. *Nat. Hum. Behav.* **4**, 1–12 (2020).
10. Webster, S. W. & Abramowitz, A. I. The ideological foundations of affective polarization in the us electorate. *Am. Politics Res.* **45**, 621–647 (2017).
11. Nordbrandt, M. Affective polarization in crosscutting communication networks: Offline and online evidence from spain. *Front. Political Sci.* **4**, 921188 (2022).
12. Törnberg, P., Andersson, C., Lindgren, K. & Banisch, S. Modeling the emergence of affective polarization in the social media society. *PLoS ONE* **16**, e0258259 (2021).
13. Brady, W. J., Wills, J. A., Jost, J. T., Tucker, J. A. & Van Bavel, J. J. Emotion shapes the diffusion of moralized content in social networks. *Proc. Natl Acad. Sci. USA* **114**, 7313–7318 (2017).
14. Brady, W. J., McLoughlin, K., Doan, T. N. & Crockett, M. J. How social learning amplifies moral outrage expression in online social networks. *Sci. Adv.* **7**, eabe5641 (2021).
15. Rathje, S., Van Bavel, J. J. & Van Der Linden, S. Out-group animosity drives engagement on social media. *Proc. Natl Acad. Sci. USA* **118**, e2024292118 (2021).
16. Cinelli, M., De Francisci Morales, G., Galeazzi, A., Quattrociocchi, W. & Starnini, M. The echo chamber effect on social media. *Proc. Natl Acad. Sci. USA* **118**, e2023301118 (2021).
17. Bakshy, E., Messing, S. & Adamic, L. A. Exposure to ideologically diverse news and opinion on facebook. *Science* **348**, 1130–1132 (2015).
18. Bail, C. A. et al. Exposure to opposing views on social media can increase political polarization. *Proc. Natl Acad. Sci. USA* **115**, 9216–9221 (2018).
19. Tyagi, A., Uyheng, J. & Carley, K. M. Affective polarization in online climate change discourse on twitter. In *2020 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)* (2020).
20. Yarchi, M., Baden, C. & Kligler-Vilenchik, N. Political polarization on the digital sphere: A cross-platform, over-time analysis of interactional, positional, and affective polarization on social media. *Political Commun.* **38**, 98–139 (2021).
21. Mentzer, K., Fallon, K., Prichard, J. & Yates, D. Measuring and unpacking affective polarization on twitter: The role of party and gender in the 2018 senate races. *Computer Information Systems Journal Articles* 2459–2468 (2020).
22. Chen, E., Lerman, K. & Ferrara, E. Tracking social media discourse about the covid-19 pandemic: Development of a public coronavirus twitter data set. *JMIR Public Health Surveill.* **6**, e19273 (2020).
23. Chang, R.-C., Rao, A., Zhong, Q., Wojcieszak, M. & Lerman, K. # roeoverturned: Twitter dataset on the abortion rights controversy. In *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 17, 997–1005 (2023).
24. Tangherlini, T. R. Toward a generative model of legend: Pizzas, bridges, vaccines, and witches. *Humanities* **7**, 1 (2017).
25. Shahsavari, S., Holur, P., Wang, T., Tangherlini, T. R. & Roychowdhury, V. Conspiracy in the time of corona: automatic detection of emerging covid-19 conspiracy theories in social media and the news. *J. computational Soc. Sci.* **3**, 279–317 (2020).
26. Tang, J., Liu, J., Zhang, M. & Mei, Q. Visualizing large-scale and high-dimensional data. In *Proceedings of the 25th International Conference on World Wide Web*, 287–297 (2016).
27. Iyengar, S. & Westwood, S. J. Fear and loathing across party lines: New evidence on group polarization. *Am. J. Political Sci.* **59**, 690–707 (2015).
28. Rao, A., Morstatter, F. & Lerman, K. Retweets amplify the echo chamber effect. In *Proceedings of the 2023 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, ASONAM '23*, 30–37 (2024).
29. Weng, L., Menczer, F. & Ahn, Y.-Y. Virality prediction and community structure in social networks. *Sci. Rep.* **3**, 1–6 (2013).
30. Dredze, M., Paul, M. J., Bergsma, S. & Tran, H. Carmen: A twitter geolocation system with applications to public health. In *AAAI workshop on HIAI*, vol. 23, 20–24 (2013).
31. Pennebaker, J. W., Francis, M. E. & Booth, R. J. Linguistic inquiry and word count: Liwc 2001. *Mahway: Lawrence Erlbaum Associates* **71**, 2001 (2001).
32. Mohammad, S. M. & Turney, P. D. Emotions evoked by common words and phrases: using mechanical turk to create an emotion lexicon. In *Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, 26–34 (2010).
33. Warriner, A. B., Kuperman, V. & Brysbaert, M. Norms of valence, arousal, and dominance for 13,915 english lemmas. *Behav. Res. methods* **45**, 1191–1207 (2013).
34. Chochlakis, G. et al. Leveraging label correlations in a multi-label setting: A case study in emotion. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5 (2023).
35. Mohammad, S., Bravo-Marquez, F., Salameh, M. & Kiritchenko, S. Semeval-2018 task 1: Affect in tweets. In *Proceedings of the 12th international workshop on semantic evaluation*, 1–17 (2018).
36. Barberá, P. Birds of the same feather tweet together: Bayesian ideal point estimation using twitter data. *Political Anal.* **23**, 76–91 (2015).
37. Conover, M. et al. Political polarization on twitter. In *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 5, 89–96 (2021).
38. Badawy, A., Ferrara, E. & Lerman, K. Analyzing the digital traces of political manipulation: the 2016 russian interference twitter campaign. In *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 258–265 (2018).
39. Le, H. et al. Measuring political personalization of google news search. In *The World Wide Web Conference*, 2957–2963 (2019).
40. Nikolov, D., Flammini, A. & Menczer, F. Right and left, partisanship predicts (asymmetric) vulnerability to misinformation. *Harvard Kennedy School (HKS) Misinformation Review* (2020).

41. Rao, A. et al. Political partisanship and antiscience attitudes in online discussions about covid-19: Twitter content analysis. *J. Med. Internet Res.* **23**, e26692 (2021).
42. MBFC. Media bias/fact check news. Media Bias Fact Check. Available at: <https://mediabiasfactcheck.com> (2023). Accessed: 2024-04-11.
43. Schaeffer, K. Nearly three-in-ten americans believe covid-19 was made in a lab. Pew Research Center. Available at: <https://www.pewresearch.org/short-reads/2020/04/08/nearly-three-in-ten-americans-believe-covid-19-was-made-in-a-lab/> (2020). Accessed: 2024-04-11.
44. Connaughton, A. Those on ideological right favor fewer covid-19 restrictions in most advanced economies. Pew Research Center. Available at: <https://www.pewresearch.org/lockdown-continues-across-sydney-as-nsw-continues-to-record-new-covid-19-cases/> (2021). Accessed: 2024-04-11.
45. Rojas, R. Masks become a flash point in the virus culture wars. New York Times. Available at: <https://www.nytimes.com/2020/05/03/us/coronavirus-masks-protests.html> (2020). Accessed: 2024-04-11.
46. Pierri, F. et al. Online misinformation is linked to early covid-19 vaccination hesitancy and refusal. *Sci. Rep.* **12**, 1–7 (2022).
47. Rathje, S., He, J. K., Roozenbeek, J., Van Bavel, J. J. & van der Linden, S. Social media behavior is associated with vaccine hesitancy. *PNAS Nexus* **1**, pgac207 (2022).
48. Mackenzie, C. Abortion and embodiment. *Australas. J. Philos.* **70**, 136–155 (1992).
49. Hartig, H. Wide partisan gaps in abortion attitudes, but opinions in both parties are complicated. Pew Research Center. Available at: <https://www.pewresearch.org/short-reads/2022/05/06/wide-partisan-gaps-in-abortion-attitudes-but-opinions-in-both-parties-are-complicated/> (2022). Accessed: 2023-04-29.
50. Rao, A., Guo, S., Wang, S. Y. N., Morstatter, F. & Lerman, K. Pandemic culture wars: Partisan differences in the moral language of covid-19 discussions. In *2023 IEEE International Conference on Big Data (BigData)*, 413–422 (2023).
51. Rao, A., Chang, R.-C., Zhong, Q., Lerman, K. & Wojcieszak, M. Tracking a year of polarized twitter discourse on abortion. *Preprint at: https://arxiv.org/abs/2311.16831* (2023).
52. Bollen, J., Gonçalves, B., Ruan, G. & Mao, H. Happiness is assortative in online social networks. *Artif. Life* **17**, 237–251 (2011).
53. Lerman, K. & Ghosh, R. Information contagion: An empirical study of the spread of news on digg and twitter social networks. In *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 4, 90–97 (2010).
54. Vosoughi, S., Roy, D. & Aral, S. The spread of true and false news online. *Science* **359**, 1146–1151 (2018).
55. Garimella, K., De Francisci Morales, G., Gionis, A. & Mathioudakis, M. Political discourse on social media: Echo chambers, gatekeepers, and the price of bipartisanship. In *Proceedings of the 2018 World Wide Web Conference*, 913–922 (2018).
56. Jiang, J., Ren, X. & Ferrara, E. Retweet-bert: political leaning detection using language features and information diffusion on social networks. In *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 17, 459–469 (2023).
57. Wilson, T. & Starbird, K. Cross-platform disinformation campaigns: lessons learned and next steps. *Harvard Kennedy School Misinformation Review* **1** (2020).
58. Evkoski, B., Ljubešić, N., Pelicon, A., Mozetič, I. & Kralj Novak, P. Evolution of topics and hate speech in retweet network communities. *Appl. Netw. Sci.* **6**, 1–20 (2021).

Acknowledgements

This work was supported in part by AFOSR grant FA9550-20-1-0224, DARPA under contract HR001121C0168 and NSF under grant CCF 2200256.

Author contributions

K.L. conceived the experiment, A.R. prepared the data, D.F., A.R., K.L., and Z.H. conducted the analysis and analyzed the results. All authors contributed to writing and reviewing the manuscript.

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/s44260-024-00008-w>.

Correspondence and requests for materials should be addressed to Kristina Lerman.

Reprints and permissions information is available at <http://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 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2024