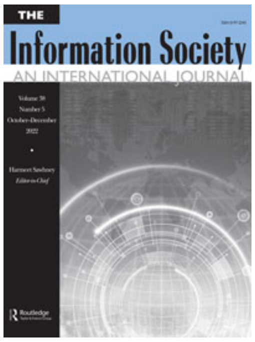

The Information Society
An International Journal

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/utis20

Social network dynamics, bots, and community-based online misinformation spread: Lessons from anti-refugee and COVID-19 misinformation cases

Lichen Zhen, Bei Yan, Jack Lipei Tang, Yuanfeixue Nan & Aimei Yang

THE INFORMATION SOCIETY
https://doi.org/10.1080/01972243.2022.2139031

Routledge
Taylor & Francis Group

Check for updates

Social network dynamics, bots, and community-based online misinformation spread: Lessons from anti-refugee and COVID-19 misinformation cases

Lichen Zhen[a], Bei Yan[b], Jack Lipei Tang[a], Yuanfeixue Nan[a] and Aimei Yang[a]

[a]Annenberg School for Communication and Journalism, University of Southern California, Los Angeles, California, USA; [b]School of Business, Stevens Institute of Technology, Hoboken, New Jersey, USA

ABSTRACT
Networked social influence and strategic information manipulation are two social mechanisms fueling misinformation spread in online communities. However, it is unclear how these two mechanisms differ in their impacts. We conducted social network analyses on two online communities sharing misinformation concerning refugees in 2016 and COVID-19 in 2020. The results robustly showed that online misinformation spread is transitive and positively associated with members' embedded authority (i.e., the extent to which members' information is exclusively shared within the focal community). At the same time, strategic misinformation sharing by members of high community loyalty (i.e., targeted information sharing within the community) is less likely to gain momentum. The impact of bots on misinformation is contingent. Findings suggest that networked social influence is a more powerful driver of misinformation spread than strategic information manipulation.

ARTICLE HISTORY
Received 6 January 2021
Accepted 15 February 2022

KEYWORDS
Bots; misinformation; online communities; social influence; social networks

During major elections and catastrophic public crises such as the COVID-19 pandemic, millions of people worldwide saw misinformation flood their social media feeds and find its way into the mass media ecosystem (Cinelli et al. 2020). Although the spread of misinformation is not always intentional, it often causes confusion, hinders decision-making, and sometimes even leads to the loss of human lives (Centers for Disease Control and Prevention 2020; Jang et al. 2018; Mocanu et al. 2015; Shin et al. 2018; World Health Organization 2020). While the free flow of accurate information has long been considered a pillar of democracy, the prevalence of misinformation has undermined traditional journalism institutions as well as posed direct threats to democracy (Hochschild and Einstein 2015). In many countries, the rise of populism paired with decreased trust in established news sources and coordinated misinformation campaigns has led to polarized public opinion and sentiment, obstructing political institutions' proper functioning, especially when they are needed the most (Lazer et al. 2018).

Recognizing the scope of the misinformation problem and its potential impact on institutions of democracy and public safety, research on online misinformation spread has surged in recent years (Guo and Vargo 2020; Wardle 2017; Weeks and Gil de Zúñiga 2021). Our literature review suggests that scholars have found that two social mechanisms fuel the dissemination of misinformation in online communities: networked social influence (Hameleers and van der Meer 2020; Vosoughi, Roy, and Aral 2018) and strategic information manipulation (Bessi and Ferrara 2016; Mejias and Vokuev 2017). Networked social influence refers to misinformation spread due to the norms and structures of social networks. Mechanisms such as transitivity, homophily, and embeddedness are found to be related to users' selective sharing and commenting behavior online. Such behaviors promote the spread of partisan information and even fake news (Shin et al. 2017; Wang and Song 2020). Strategic information manipulation to distribute misinformation is committed by, for instance, social media accounts set up by fake news sites (Vargo, Guo, and Amazeen 2018) and bots (Ferrara 2017), agenda-setting and seeding of misinformation.

However, it remains unclear how these two forces work together to drive misinformation spread in online communities and, in particular, if they have different effects on the diffusion of misinformation

CONTACT Lichen Zhen lichen.zhen@usc.edu Annenberg School for Communication and Journalism, University of Southern California, 3502 Watt Way, Los Angeles, CA 90007, USA.
Supplemental data for this article can be accessed online at https://doi.org/10.1080/01972243.2022.2139031

online. This study thus seeks to examine how social network dynamics and strategic actions shape misinformation transmission in online communities and compares the impacts of the two types of processes. We view online communities as networks of social media users connected by information sharing, conversations, or other forms of communication.

Our study examines two Twitter-based online communities that shared misinformation in two different cases: an online community that spread misinformation about refugees in 2016 and another community that disseminated misinformation about the origins of COVID-19 in early 2020. Social network analyses of both communities showed that online misinformation sharing networks are highly transitive, suggesting an association between network closure and misinformation spread. Misinformation shared by members with high embedded authority, or those whose information is exclusively shared by the online community members, is more likely to disseminate. By contrast, misinformation shared by bots as well as members of high community loyalty who target specific communities is less likely to circulate. The consistent findings from two cases suggest that although strategic information manipulations are conducted to disseminate misinformation online, networked social influence is a more powerful force driving misinformation spread in online communities. Recommendations for policy and practice are provided according to our findings.

Literature review

Misinformation is information that incorporates inaccurate or distorted content that was shared within misleading contexts (Weeks and Gil de Zúñiga 2021). Given the challenges posed by misinformation, a growing number of scholars have devoted considerable attention to this issue (Shin et al. 2017; Southwell and Thorson 2015; Weeks and Gil de Zúñiga 2021). Many scholars have investigated factors that impact the cognitive processing of misinformation at the individual level. For example, they have examined the role of visual exemplars (Dixon et al. 2015), issue-specific knowledge (Krishna 2017), emotions (Weeks 2015), social media functionality, and experts opinion (Vraga and Bode 2017) in affecting the effort to correct misinformation. Although this stream of research yields important insights, it does not reveal why and how misinformation diffuses among people in social groups (Pasek, Sood, and Krosnick 2015; Weeks 2015).

In this study, we follow research studying how social dynamics in online communities contribute to misinformation spread (Bennett and Livingston 2018; Guo and Vargo 2020; Southwell and Thorson 2015). This line of research examines the social processes and structures that drive misinformation dissemination and takes mechanisms underlying the spread of misinformation to be embedded in a broader social, political, and technological context (Shin et al. 2017; Shin et al. 2018).

Our review of the literature suggests that two major processes are driving community-based online misinformation spread: (1) networked social influence and (2) strategic information manipulation. Networked social influence comes from the social structures of the online communities that people are embedded in. Research has suggested that networked communities formed by people with similar political ideologies on Twitter disseminated political misinformation among community members (Guo, Rohde, and Wu 2020; Shin et al. 2017; Shin and Thorson 2017; Wang and Song 2020). Influential users of the online communities had strong impacts on other members and were major contributors to the diffusion of misinformation in these networks (Guo, Rohde, and Wu 2020; Shin et al. 2018; Vosoughi, Roy, and Aral 2018). The influence of homophilous social networks is potent because individuals often select information consistent with their political beliefs (Hameleers and van der Meer 2020; Waisbord 2018).

The second mechanism which has perhaps drawn more popular and academic attention is strategic (mis) information manipulation. Web 2.0 has ended the monopoly of mass media on information production and distribution and allowed citizens to "generate, consume, or distribute" information (Mejias and Vokuev 2017, 1027; see also Bennett and Livingston 2018). While technological advancement has enabled more democratized information production, it also enabled strategic actors to influence online sphere. Scholars have observed coordinated movements to propagate misinformation and influence public opinions committed by fake news sites (Bennett and Livingston 2018; Vargo, Guo, and Amazeen 2018; Weeks and Gil de Zúñiga 2021) and social media bots (Del Vicario et al. 2016; Wojcik et al. 2018). A recent study conducted by researchers from Brown University found that 25% of tweets on topics related to climate change were posted by bots (Milman 2020), which mainly spread denials of global warming or rejections of climate science. Researchers have also detected significant bot activities on Twitter during the 2016 U.S. presidential election (Bessi and Ferrara 2016) and the 2017 French election (Ferrara 2017).

Nevertheless, research has yet to delineate or compare the effects of networked influence and strategic manipulation in online communities. More research is needed to explicate how these two forces function to fuel online misinformation spread. We conceptually distinguish the effects of networked social influence and strategic information manipulation on disseminating misinformation online and compare the impacts of the two processes. In the next section, we turn our attention to literature from online communities, social networks, and bots and identify social and strategic processes in online communities relevant to the spread of misinformation.

The mechanisms of misinformation spread in online communities

The rise of the Internet has enabled users to build social networks with others who share similar values and goals without forming formal organizations or being limited by spatial constraints (Willson 2010). Online communities are such densely connected social networks (Wang, Yang, and Thorson 2021). Although early research conceptualized the Internet as a public sphere that enables meaningful debates among diverse participants over public issues, recent research has painted a rather different picture. Scholars have observed that online spaces are highly polarized and found that individual users prefer to interact with like-minded others (Chan and Fu 2017; Dvir-Gvirsman 2017). These features make online communities susceptible to both networked social influence and strategic information manipulation.

Networked community interactions

Networked communities tend to exhibit common relationship formation tendencies such as transitivity – two members are sharing a conversation, and one of them introduces the other to her friend, enabling their friends to become friends (Holland and Leinhardt 1971; Monge and Contractor 2003; Wang, Yang, and Thorson 2021).

Transitivity in social networks is marked by an additional path between two individuals through a third party. Intransitive relationships, in contrast, refer to the social phenomenon wherein two individuals are both connected to an intermediary but do not form a direct tie. Transitivity is a common pattern in social groups that indicates balanced social relationships (Burt 2005) and is associated with increased cohesion and trust within groups (Monge and Contractor 2003). Intransitive ties indicate unbalanced relationships among individuals because the person at the intermediary position can control the flow of information between other individuals. This unbalanced structure may cause discomfort and become a source of distress among individuals if the intermediary intentionally hides or distorts information (Doran, Alhazmi, and Gokhale 2013).

Our review of previous research suggests that misinformation spreading in online communities is likely to be transitive. First, in close-knit clusters such as online communities, two members who shared misinformation from the same person may find each other more credible or trustworthy, and thus are more likely to disseminate one another's information and close the triad (Burt 2005). Block (2015) observes that the tendency toward transitivity is so strong in social networks that it is only offset by hierarchical relationships. Since online communities often lack hierarchy (Willson 2010), we expect there to be a strong inclination in online communities to eliminate the intermediary and connect the two unconnected individuals in an intransitive triad (Doran, Alhazmi, and Gokhale 2013). Second, individuals are more likely to find repetitive information to be credible (e.g., several close contacts saying the same thing), especially when people are trying to make sense of new information or situations with high levels of uncertainty (Koch and Zerback 2013). Research has shown that repeating a message makes the message makes it seem more credible than communicating it once (Dechêne et al. 2010). Similarly, sharing similar misinformation from multiple sources in a community may help enhance the perceived credibility (Watts 2002). We therefore reason that individuals are more likely to disseminate the message from another member if the two are connected through a shared third party via their misinformation sharing activities, thus forming transitive triads in their information sharing network. Finally, although information diffusion research has found that information diffusion in large scale networks may form a cascading structure with low connectivity and low transitivity (Iribarren and Moro 2011), research examining the diffusion of misinformation within more tightly connected networks has found that misinformation diffused in clusters often show strong tendencies of transitivity (Lai and Wong 2002). Building on previous research on transitivity, we propose:

> H1: Misinformation spread in communities is more likely to be transitive than by chance alone.

Political homophily

Political homophily refers to the phenomenon that information sharing is more likely to occur among people who share similar political ideologies

– systematic political views held by individuals, groups, and cultures (Thorson 2016). In online communities of politically like-minded members this can lead to the "echo chamber effect" (Colleoni, Rozza, and Arvidsson 2014). It occurs when clusters of online community members interact primarily with others holding similar political ideologies and avoid political opposition and uncomfortable discussions. Brainard (2009, 598) observed the tendency that "people seek out only like-minded others and thereby close themselves off from ideological opposition, alternative understandings, and uncomfortable discussions." Political homophily can be explained by cognitive dissonance theory, which suggests that individuals seek to avoid discomfort caused by inconsistent political attitudes in information exposure and therefore generally prefer political information that aligns with their ideological attitudes (Song and Boomgaarden 2017).

Network research has generally found that tie formation occurs more often among people who demonstrate similar behaviors or are like-minded (McPherson, Smith-Lovin, and Cook 2001). Furthermore, recent studies have identified a direct connection between political homophily and misinformation sharing (Shin et al. 2017; Weeks and Gil de Zúñiga 2021). For example, Shin et al. (2017) analyzed the 2012 U.S. presidential election campaign Twitter data and found that the network structure of political rumor-tellers was based on partisanship. In other words, misinformation is more likely to be shared among politically homophilous actors. On Twitter, when users show certain political tendencies through their self-provided profile information or the content of tweets, they may attract interaction of others who have similar profiles/views, including the sharing of misinformation. Hence, we propose the following hypothesis:

> H2: Misinformation spreading is more likely to occur among users sharing similar political ideologies.

Cyberbalkanization

So far, we have discussed how social media users tend to interact with like-minded others as pairs or small clusters. At the community level, such interaction patterns can aggregate and create a phenomenon termed cyberbalkanization – members of online communities primarily interacting with other members at the cost of interactions with nonmembers (Chan and Fu 2017). Previous research measured the social embeddedness of online community members using the cyberbalkanization index, or the degree to which users interact more with others within their community than outside of their community. The more a user exclusively interacts with others within her community, the higher her cyberbalkanization index (Chan and Fu 2017). Our study builds upon previous research and differentiates member's level of cyberbalkanization based on the direction of information sharing. This is because, as further explained below, variation in information exchange activities may be suggestive of either networked influence or strategic manipulation.

Embedded authority

Embedded authority captures the extent to which members' information is shared within the focal community compared to outside of the community. Members who have high embedded authority have high scores on cyberbalkanization indices in terms of who shares their information. The information they spread is most likely to reach and influence members in their communities but not non-community members. Embedded authority does not equal absolute popularity within a focal community. It is a relative concept comparing one's information influence within versus outside of the community. One can have high embedded authority if they have moderate shares within the community and no shares beyond the community. Similarly, one can have low embedded authority if they are popular and widely shared across different communities.

High embedded authority grants members the power to influence their peers in their focal communities who are consequently, more likely to drive online misinformation spread. Individuals tend to place more value on information disseminated by members who exclusively participate in their communities (Chan and Fu 2017; DeMarzo, Vayanos, and Zwiebel 2003) due to selective exposure, social comparison, and social corroboration (Sunstein 2007). Selective exposure occurs as community members selectively read and share information from like-minded peers (Holbert, Garrett, and Gleason 2010). Social comparison is the process by which group members adjust their view toward the well-received information shared by their in-group members. Social corroboration refers to community members reconfirming the information with other members, thus becoming more confident and reinforced in their opinions. Because misinformation shared by members with high embedded authority is solely shared within communities, consumption of such information distinguishes members from nonmembers. Correspondingly, misinformation is more likely to be considered as in-group information accepted by the community members. As a result, community members will be more likely to discuss, disseminate, and confirm the misinformation shared by embedded authorities, which reinforces the influence of embedded authorities on them. Therefore, we hypothesize that:

> H3a: Embedded authority is positively associated with the likelihood that a member's misinformation is shared within the community.

Community loyalty

Unlike the embedded authority, members' community loyalty relates to the level of information sharing within a community compared to outside of the community. Members with high community loyalty target their information sharing to a particular community. These members score high on the cyberbalkanization index in terms of information sharing activities, meaning that they focus on sharing information to specific communities but not others. Like embedded authority, community loyalty is a relative construct comparing members' information sharing within over that outside of a community. Members can be high in community loyalty if they share a moderate amount of information exclusively in one community or can have low community loyalty if they share information across different communities.

Although the cyberbalkanization index did not distinguish cyberbalkanization in terms of embedded authority and community loyalty (Chan and Fu 2017), our study differentiates the two constructs. The study proposes that they indicate different actors' social behavior and relate to online misinformation spread through different social processes. As discussed above, members with high embedded authority are mainly shared within communities, but not others. Since members cannot directly manipulate others' behavior, embedded authority is more likely to be gradually established through their community's social connections and interaction processes.

By contrast, members with high levels of community loyalty can be strategic outsiders who target online communities to influence their views. Previous research has reported that partisan media outlets and fake news sites focus their misinformation campaigns on groups of homogeneous audiences to fuel the spread of online misinformation (Dvir-Gvirsman 2017; Vargo, Guo, and Amazeen 2018; Weeks and Gil de Zúñiga 2021). Targeted information sharing by strategic actors may further amplify the influence of misinformation due to community members' in-group bias that favor their cyberbalkanized members, as discussed above. In other words, because of group identity and in-group biases, if community members see someone exclusively sharing information to their community, they are more likely to further relay and discuss such information.

> H3b: Community loyalty is positively associated with the likelihood that a member's misinformation is shared within the community.

However, unlike those with high embedded authority, members high in community loyalty occupy a different network position within their communities. Their strategic purpose of circulating information to the community may be more easily identifiable to the community members. As such, although we also expect members' with high community loyalty have a critical role in the online misinformation spreading process, we are also interested in examining if the two types of social behaviors influence misinformation sharing in online communities differently.

> RQ1: Do embedded authority and community loyalty impact misinformation spread in online communities differently?

Bots

Another type of strategic actor in online communities is the social media bot. Bots are often designed to simulate human activities on social media by generating content and interacting with other users. They play a critical role in misinformation spread. Pew has reported that as high as 66% of the links to popular websites on Twitter were shared by bots (Wojcik et al. 2018). Whereas some social media bot activity is benevolent, plenty of research has detected coordinated bot activities to spread misinformation and influence public opinions (Bessi and Ferrara 2016; Del Vicario et al. 2016; Ferrara 2017; Milman 2020). However, the specific effects of bots on the magnitude of online misinformation sharing remain ambiguous. Given the prevalence of bot activities on Twitter and its relevance to misinformation spread, we are interested in exploring the role of bots in the diffusion of misinformation online. Therefore, we propose the following research question:

> RQ2: Are accounts classified as bots more likely to drive misinformation spread in online communities than non-bot accounts?

Methodology

In this study, we examined two cases of misinformation spread in two online communities on Twitter.

The first one is related to a story published by Judicial Watch about refugees in June 2016. The second one is about falsified claims on Twitter that COVID-19 was a bioengineered virus with HIV insertion.

In the first case (hereafter the anti-refugee misinformation case), Judicial Watch, a conservative advocacy group, claimed that an undocumented Middle Eastern woman and "Islamic refugee" was arrested in Luna County, New Mexico and that she was carrying regional gas pipeline plans. Fact-checking website snopes.com contacted both Luna County Sheriff's department and the U.S. Border patrol to verify the news, and both offices denied the validity of this story. Despite little evidence to support the claims in this news story, the story was retweeted by thousands of unique Twitter users.

In the second case (hereafter the COVID-19 misinformation case), a team of researchers published a paper on the preprint biological website bioRxiv.org, which claimed to have found "uncanny similarities between the amino acid structure in SARS-CoV-2 and HIV" and that this is unlikely to be found in nature (Kasprak 2020). Although experts have confirmed that this similarity between COVID-19 and HIV is due to random chance (Kasprak 2020), the abovementioned paper triggered a widespread rumor on Twitter. Some Twitter users used this piece of misinformation as evidence to support a conspiracy theory that COVID-19 was a bioengineered virus.

Data

Data for the two cases were collected from two different sources. For the first case, we used DiscoverText, a Twitter analytics tool (Shulman 2011), to purchase tweets with "refugee" as the search keyword in the news. This dataset was purchased in 2016. For the second, COVID-19 misinformation case, we utilized a large public Twitter database collected through a comprehensive set of COVID-19 related keywords (Chen, Lerman, and Ferrara 2020). Both datasets included both tweet-level variables such as the full text of the tweets and user-level variables such as users' bios.

Our sampling periods (see Table 1) were chosen because they covered the time window when the abovementioned misinformation was highly retweeted. We chose the date for the time when the misinformation first appeared as the start date of the selection window. For the anti-refugee misinformation case, we did not find tweets containing the three keywords after July 1 in the purchased dataset and thus chose July 1 as the end date. For the COVID-19 case, the peak of the misinformation spread was between January 31 and February 4. We chose February 29 as the end date because it was a week after the last spike of the misinformation spread.

Table 1. Network descriptive statistics.

	2016 Refugee case	2020 COVID19 case
Data collection date	June 15th–July 1st, 2016	January 31st–February 4, 2020
Total tweets	488,401	22,560,356
The complete network		
Nodes	172,675	4,718,959
Ties	292,740	15,842,975
Density	0.00000982	0.00000071
The core misinformation spread network		
Nodes	522	308
Ties	806	420
Density	0.002964	0.004442

Case 1: Refugee information network

We obtained 488,401 English tweets mentioning the term "refugee" through DiscoverText. Defining retweeting as a network tie, we constructed the *complete* refugee information network using these tweets. This network contained 172,675 unique users (nodes) and 292,740 directed retweeting relationships (ties). The descriptive statistics of the network are illustrated in Table 1.

To identify a retweet network circulating the misinformation regarding the arrest of an Islamic refugee, we first used three case-insensitive search terms ("Mexico," "gas," and "Islamic refugee") and the "AND" operator to identify tweets that mentioned all of the keywords discussed above. This search located 5,507 tweets mentioning all three words, among which 3,278 were retweets. We constructed a misinformation network based on the search results, and this network had 3,070 nodes (Network density = 0.00035). This misinformation network was over 30 times more densely connected than the complete network, suggesting that it is a more close-knitted community compared to the complete network.

To ensure we looked at a community of frequent interacting users, we removed isolates (users who did not retweet or were retweeted by others) and pendants (users who only connect to the network through one tie). As a result, we obtained the *core* misinformation spread network of 522 nodes connected by 806 ties (Network density = .002964). An additional manual check was conducted on the tweets sent by these 522 users to make sure that the tweets were about the misinformation published by Judicial Watch.

Case 2: COVID information network

We obtained 22,560,356 English tweets between January 22 and February 29, 2020 from the COVID-19 tweet dataset (Chen, Lerman, and Ferrara 2020). Among these tweets, 15,842,975 were retweets. We constructed the *complete* COVID-19 information network based on these retweets. The complete network had 4,718,959 nodes and 15,842,975 ties.

We used the same approach as case one to identify the *core* misinformation spread network in the complete COVID-19 information network. We first extracted a misinformation retweet network spreading the claims that COVID-19 was bioengineered. We used "AND" operator with two case-insensitive keywords "bioweapon" and "HIV" to search for all tweets within Chen, Lerman, and Ferrara (2020)'s dataset. 5,885 English tweets were obtained from the search. Additional manual checking revealed that 15 tweets included counter-misinformation messages. We removed these 15 tweets, leaving 5,870 tweets containing the misinformation. Using these tweets, we built a misinformation retweet network that consisted of 4,866 users and 4,784 ties. The density of the network is 0.00020. This misinformation retweet network was over 280 times more densely connected than the complete network, indicating a tightly knitted misinformation-spreading community compared to the complete network. After removing isolates and pendants, we obtained the core misinformation spread network of 308 nodes connected by 420 ties (Network density = .004442). The COVID-19 misinformation community was a smaller and slightly denser network than the anti-refugee misinformation community.

Figures 1 and 2 provide visualizations of the retweet network within the two core misinformation spread networks. Nodes with the highest betweenness centralities were labeled in the figures. We can observe that Judicial Watch is the most central node for the core misinformation network in case one, and TrumpGAGirl is the node with the highest betweenness centrality in case two.

Independent variables

Our model included four independent variables and three control variables. The descriptive statistics of the variables are summarized in Table 2. The first independent variable political ideology and one control variable Judeo-Christian identity were assessed by human coding of users' Twitter bios. The user bio information was collected in 2016 for the first case and 2020 for the second case. Three coders independently coded the bios. Each account was coded twice by two different coders. Eighty-seven accounts in the anti-refugee core misinformation network and forty-eight in the COVID-19 core misinformation network were not coded for political ideology and Judeo-Christian identity because these users did not provide a bio at the time of data collection.

Political ideology

Political ideology was coded into three categories: liberal, conservative, and no apparent political inclination. Users were coded as liberal when they identified themselves as a Democrat, advocated progressive policies, or supported democratic candidates in their bios. Users were coded as conservative if they identified themselves as a Republican, advocated conservative policies, or supported Republican candidates. Bios that did not mention partisan-leaning content were categorized as having no apparent political inclination. Although there is a possibility that a user might disguise the actual political ideology by faking the bio, we argue that it is the self-presentation of political ideology that matters in information spread on social media. An experiment study (Lee, Kim, and Coe 2018) found that self-presented political ideology in the bio influences others' perception and communication behavior regardless of user partisanship. Users perceived the news shared by others self-identified as from the opposing party as more biased.

For the first misinformation case, the coders agreed on 89% of the accounts coded in the initial coding, reaching a Cohen's kappa of 0.68. The coders then discussed the coding and independently coded a sample of users (20%) again with a kappa of 0.80. For the second misinformation case, intercoder reliability reached a Cohen's kappa of 0.77. The inter-coder reliability is acceptable based on prior research (Lombard, Snyder-Duch, and Bracken 2002). For both cases, if an account was coded differently, the three coders together reread accounts' bios and discussed until a full agreement was reached. As shown in Table 2, the majority of the members in both communities demonstrated a conservative political tendency. This observation is consistent with prior research showing that conservatives are much more likely to be the targeted audience of misinformation spreaders and to believe or share misinformation (Hjorth and Adler-Nissen 2019; Freelon and Wells 2020).

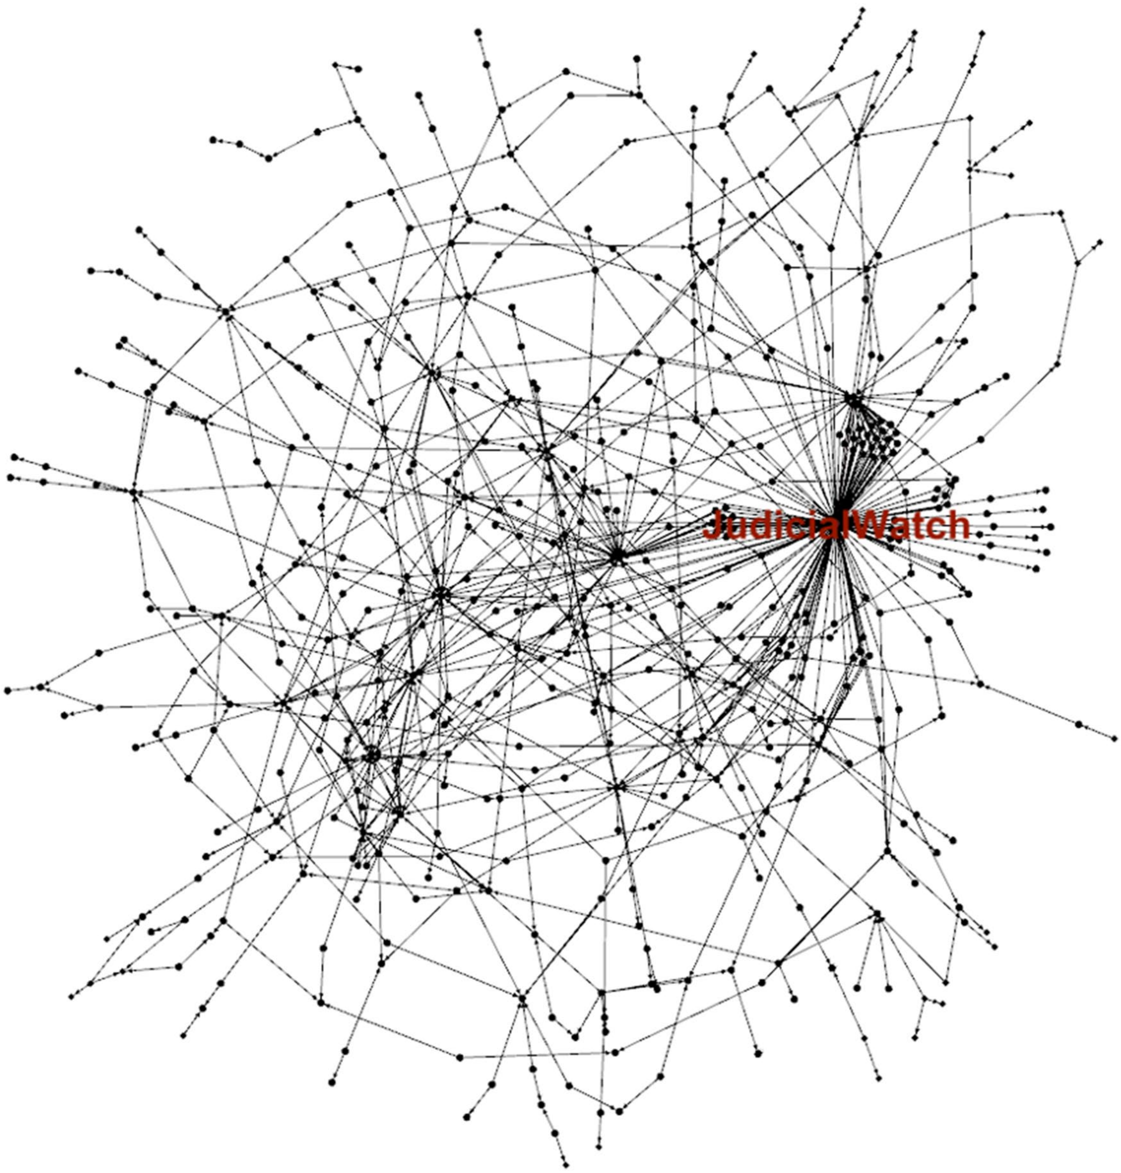

Figure 1. Network visualization of the retweet network for the refugee misinformation community.
Note. Judicial Watch is the central node with the highest betweenness centrality.

Cyberbalkanization

The cyberbalkanization index is calculated as the proportion of users' information exchange within an online group to their total information exchange activities (Chan and Fu 2017). As mentioned before, we differentiated two aspects of cyberbalkinzation based on users' information sharing behaviors – the level of *embedded authority* and *community loyalty*.

- *Embedded authority* represents the extent to which members' messages are mainly retweeted by members in the community as compared to outside of the community. This index was calculated as the number of times a member in the core misinformation spread network was retweeted within the core network divided by their ties in the complete information-sharing network (including their retweets and being retweeted by others). A high embedded authority index suggests that a community member was mainly retweeted by members in the core misinformation spread network, but not users in the complete information network.
- *Community loyalty* captures the extent to which a member in the core misinformation spread network shares messages within versus outside of the community. It was operationalized as the number of times a member retweeted information to the core misinformation spread network divided by their total ties in the complete

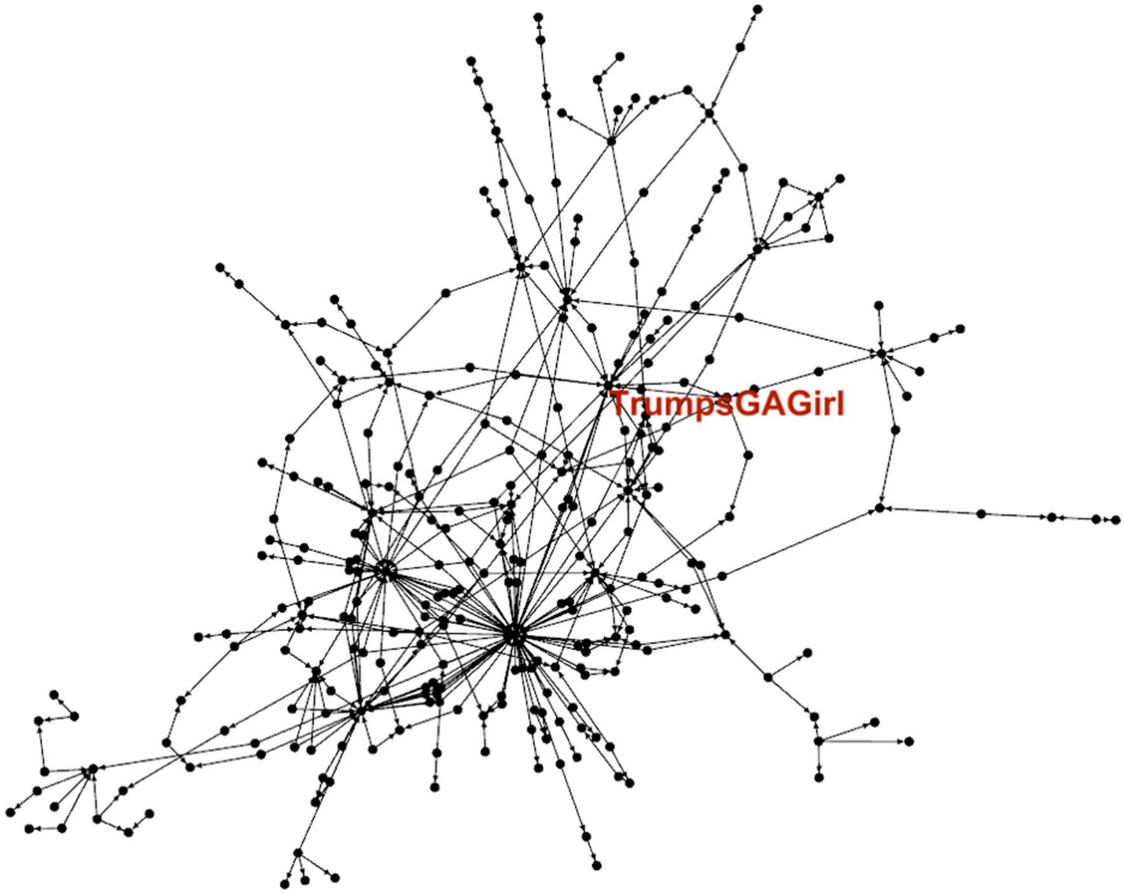

Figure 2. Network visualization of the retweet network for the COVID-19 misinformation community.
Note. TrumpGAGirl is the central node with the highest betweenness centrality.

Table 2. Descriptive statistics of the core misinformation spread networks.

	2016 Anti-refugee case	2020 COVID19 case
Political ideology		
Liberal	3	0
Conservative	368	103
Nonpartisan	64	157
NA	87	48
Cyberbalkanization	*M* (*SD*)	*M* (*SD*)
Embedded authority	0.07 (0.15)	0.02 (0.07)
Community loyalty	0.13 (0.18)	0.04 (0.10)
Bot		
Yes	140	91
No	382	214
NA	0	3
	M (*SD*)	*M* (*SD*)
In-degree in the complete information network	22.17 (94.05)	265.20 (994.34)
Out-degree in the complete information network	13.95 (15.24)	92.17 (153.16)
Judeo-Christian identity		
Yes	131	224
No	304	36
NA	87	48

information-sharing network. The higher the ratio, the more likely a user only shared information within the core misinformation network.

The graphic illustrations of nodes having high versus low levels of embedded authority and community loyalty can be found in Figure 3. In all four graphs, the black node in the middle represents the ego. The shapes and shades indicate the communities each node belongs to. Figure 3a and b represent an ego with high versus low embedded authority, respectively. The ego in Figure 3a has an embedded authority score of 0.8. It has four incoming ties within the community and only one tie outside of the community, indicating most of its information was shared by the community members. Quite the opposite, the ego in Figure 3b only has one incoming tie within the community but four ties initiated by different outside community members. It shows that members from different communities share the ego's information, and thus the ego has a lower embedded authority (0.2). Figure 3c and d show an ego with high versus low community

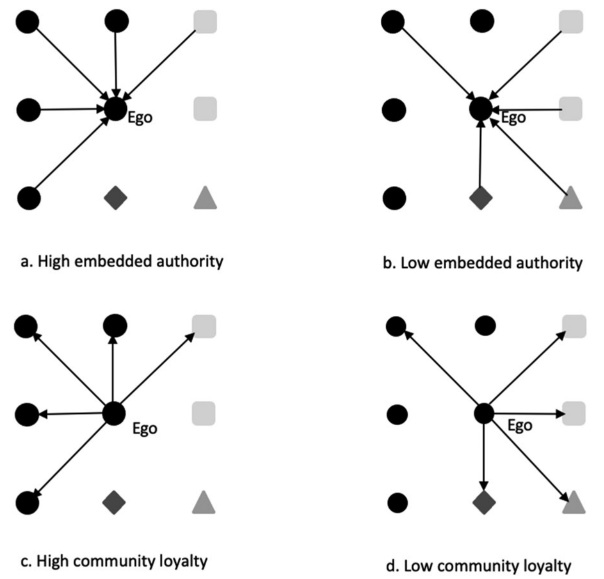

Figure 3. Egos with high and low embedded authority and community loyalty.

loyalty. The ego in Figure 3c has a community loyalty of 0.8 because it has four outgoing ties within the community and only one tie going outside of the community. Figure 3c shows that the node's information sharing activities are observed mostly within one community. By contrast, the ego in Figure 3d has only one outgoing tie within the community but four ties toward several outside communities. Therefore, its information-sharing behavior is distributed across different communities and has a relatively lower community loyalty of 0.2.

Bots

To detect bot accounts, we applied the bot detection algorithm Bot-hunter (Beskow and Carley 2018). The method analyzes Twitter accounts' features such as tweeting activities, follower and followee network, and semantics of texts applying machine learning algorithms (i.e., Random Forest model). The detection method was shown to be highly accurate in prior applications, reaching an area under the curve (AUC) of 0.994 (Beskow and Carley 2018; Magelinski, Beskow, and Carley 2019). Different from other commonly used bot detection methods (e.g., Botometer) that analyze account data from the current Twitter API, Bot-hunter can be applied to historical data already collected by the researchers. Therefore, it is a suitable algorithm for our dataset. Bot-hunter produces a probability of an account being bot. We used 0.8 as a threshold, classifying accounts receiving a Bot-hunter probability higher than 80% as bots. In the anti-refugee misinformation community, 140 (26.82%) members were classified as bots. In the COVID-19 misinformation community, 91 (29.55%) members were classified as bots. These numbers are close and consistent with the percentage of bots reported by prior research on Twitter (Milman 2020). To further check the robustness of the bot detection, four authors randomly checked a total of 100 existing Twitter accounts from each case. They labeled the account as a bot account or not based on their profiles and recent tweeting activities. The agreement between the human coding and algorithmic detection was 84% for the case 1 and 79% for case 2. While bot detection remains a challenging problem for Twitter and computer scientists and is still a matter of debate and development (Roth and Pickles 2020), this agreement shows moderately high robustness of Bot-hunter even when applied in different contexts.

Control variables

In-degree and out-degree in the complete information network

Considering that users had differing activity levels on Twitter, we controlled for users' in-degree and out-degree in the complete information network in each case. In-degree counts the number of times a focal user's tweets were retweeted by other users. Out-degree counts the number of times the focal user retweeted other users' tweets in the complete information network.

Judeo-Christian identity

In the coding process, we noticed that many users in both cases frequently used Judeo-Christian terms to identify themselves. Religious identity is tightly related to people's political ideology and information consumption (Jost et al. 2018). To control for the effect of users' representation of their Judeo-Christian values, we manually coded users who represented themselves as believers of Judeo-Christian values in their bios (for instance, users who used terms such as "#Christian," "I follow the God of Israel," and "Catholic" at least once) as Judeo-Christian and those who did not as non-Judeo-Christian. Intercoder reliability, measured by Cohen's kappa, was higher than 0.8 for both datasets. If an account was coded differently by the two independent coders, the coders discussed the bios again until an agreement was reached.

Analytical approach

Six Exponential Random Graph Models (ERGMs) were administered to test our hypotheses (Table 3). ERGMs are a class of probability models that permits statistical inferences about network configuration applying

Table 3. ERGM models predicting tie formation in misinformation spread networks.

	Anti-refugee case						COVID-19 case					
	Model 1		Model 2		Model 3		Model 4		Model 5		Model 6	
Structural effects												
Density	−5.16***	0.09	−4.06***	0.11	−4.72***	0.13	−4.06***	0.09	−0.40***	0.11	−4.00***	0.10
Cyclical Transitivity	0.38*	0.19	1.9***	0.16	0.45*	0.18	−0.07***	0.02	−0.08	0.68	−0.10***	0.02
Intransitive triads	−0.65***	0.04	−0.73***	0.04	−0.64***	0.04	−2.05***	0.19	−2.06***	0.20	−2.06***	0.20
Actor Attributes												
Homophily: Political Ideology	0.14+	0.08	0.23**	0.08	0.10	0.08	0.22*	0.10	0.22*	0.11	0.23*	0.10
Cyberbalkanization: Embedded Authority	1.21***	0.16			0.74***	0.2	0.71***	0.02			0.64***	0.02
Cyberbalkanization: Community Loyalty			−1.76***	0.20	−0.52*	0.21			−0.81+	0.44	−0.90***	0.02
Bot	0.23***	0.07	0.22***	0.06	0.23***	0.07	−0.18*	0.09	−0.20*	0.09	−0.21*	0.10
Control variables												
In-degree in the complete information network	0.00***	0.00			0.00***	0.00	−0.00*	0.00	0.00		−0.00*	0.00
Out-degree in the complete information network			−0.00***	0.00	−0.01***	0.00			0.00	0.00	0.00	0.00
Judeo-Christian identity	−0.83***	0.08	−0.49***	0.07	−0.75***	0.08	0.15	0.12	−0.15	0.13	−0.17	0.12
AIC	9328		10023		9306		4626		4629		4624	
BIC	9412		10107		9412		4702		4704		4719	

Note. $^{+}p<0.1$, $^{*}p<.05$, $^{**}p<.01$, $^{***}p<.001$.

simulation techniques (Robins et al. 2007; Wang et al. 2013). It allows hypothesis testing of specific parameter effects on network tie formation. The first three ERGMs, Model 1 to Model 3, tested hypotheses using the core anti-refugee misinformation-spreading network. The other three ERGMs, Model 4 to Model 6, tested hypotheses using the core COVID-19 misinformation-spreading network. Models 1 and 4 investigated the effects of members' embedded authority on misinformation sharing in online communities, whereas Models 2 and 5 assessed the impact of members' community loyalty. To examine whether embedded authority and community loyalty are empirically distinct, we included both measures in Model 3 and Model 6. The goodness-of fit diagnostics for the two full models (Model 3 and Model 6) are reported in the online appendix.

Results

Table 3 summarizes the results of the ERGM analyses. H1 predicted that users engaged in the misinformation spread community are more likely to have transitive ties. Since transitivity in directed networks may take various forms, the hypothesis was tested using two structural terms in the ERGM models: the gwesp and intransitive. Gwesp tests the extent to which pairs of connected nodes (i → j) are also connected through a third node (i → k → j), thus it is often used as an indicator of cyclical transitivity. Intransitivity examines the number of triads that are not closed in the network. Therefore, it can further examine the existence of transitive triads that take a different form (i → j, i → k, k → j).

Cyclical transitivity (gwesp) in the core misinformation spread network was significant and positive in case one: Model 1 (Estimates = 0.38, p < .05), Model 2 (Estimates = 1.90, p < .001) and the full Model 3 (Estimates = 0.45, p < .0.05), suggesting that cyclical transitivity in the network occurred more frequently than random in the core anti-refugee misinformation-spreading network. However, cyclical transitivity was found to be negatively significant in two models in case two: Model 4 (Estimates = −0.07, p < .001) and the full Model 6 (Estimates = −0.10, p < .001), indicating that cyclical transitivity in the network occurred less frequently than random in the core COVID-19 misinformation-spreading network.

The intransitive triad terms in the core misinformation spread network were significant and negative across all models in both cases: Model 1 (Estimates = −0.65, p < .001), Model 2 (Estimates = −0.73, p < .001), Model 3 (Estimates = −0.64, p < .0.001), Model 4 (Estimates = −2.05, p < .001), Model 5 (Estimates = −2.06, p < .001), and Model 6

(Estimates = −2.06, $p < .0.001$). This suggests that in general, fewer intransitive triads exist in the two online communities than by chance. Taken together, while less cyclical transitivity was observed in the COVID-19 network, the total number of transitive triads was high in both cases. Therefore, H1 was supported.

H2 proposed that misinformation spread is more likely to occur among users with similar political ideology. In case one, we only observed a significant and positive effect of political homophily in Model 2 (Estimates = 0.23, $p < .001$). However, this hypothesis was strongly supported in case two, as the homophily effect of political ideology was significant and positive in three models (Model 4: Estimates = 0.22, $p < .05$; Model 5: Estimates = 0.22, $p < .05$, Model 6: Estimates = 0.23, $p < .05$). Thus, H2 is partially supported.

H3 postulated that members with a high cyberbalkanization index were drivers of misinformation spread in the online communities. Models 1 and 4 tested the effect of embedded authority. Both models showed that members with high embedded authority scores are associated with more ties in the core misinformation spread network (Model 1: Estimates = 1.21, $p < .001$; Model 4: Estimates = 0.71, $p < .001$). H3a was thus supported.

H3b also predicted a facilitating role of users with high community loyalty scores during misinformation spread online. However, Models 2 and 5, which tested the influence of members' community loyalty, found an opposite, negative effect (Model 2: Estimates = −1.76, $p < .001$; Model 5: Estimates = −0.81, $p < .01$). Therefore, H3b was rejected.

In the full model (Model 3 and Model 6), we included both embedded authority and community loyalty. The results consistently showed that members with a high embedded authority score are more likely to form ties (Model 3: Estimates = 0.74, $p < .001$; Model 6: Estimates = 0.64, $p < .001$). On the other hand, nodes with a high community loyalty score were found to less likely to form ties in the core misinformation community (Model 3: Estimates = −0.52, $p < .05$; Model 6: Estimates = −0.90, $p < .001$). The results answered RQ1 and revealed that embedded authority and community loyalty played different roles in online misinformation diffusion. While members of high embedded authority were the major actors in misinformation spread, misinformation shared by those with high community loyalty was less likely to circulate.

To answer our RQ2, we analyzed whether bots were related to misinformation spread in the online community in the ERGM models. The results showed that bot accounts were more likely to spread misinformation in case one (Model 1: Estimates = 0.23, $p < .001$; Model 2: Estimates = 0.22, $p < .001$, Model 3: Estimates = 0.23, $p < .001$). In case two, the coefficients for bots showed negative significance in the models, suggesting that bots were less likely than non-bots accounts to disseminate misinformation (Model 4: Estimates = −0.18, $p < .05$, Model 5: Estimates = −0.20, $p < .05$, Model 6: Estimates = −0.21, $p < .05$). The results suggested that bots had a complicated role in misinformation spread. Given different context, they could play a either an active role or a less noticeable role compared to non-bot accounts in the misinformation spreading process.

Discussion

Misinformation is a common social problem in contemporary societies, and a growing number of studies and public policies have focused on this topic (Freelon and Wells 2020; Southwell and Thorson 2015; Vargo, Guo, and Amazeen 2018). The current study makes several unique contributions to the scholarship by elucidating how networked social influence and strategic information manipulation impact misinformation spread in online communities. Overall, our results from two cases robustly showed that misinformation spread at the community level is associated with networked social influence (i.e., transitivity, ideological homophily, embedded authority) and strategic information manipulation (i.e., community loyalty, bots). More importantly, networked social influence is more strongly associated with misinformation spread than strategic information manipulation in online communities. These findings inform our understanding of the complexity of online community-based misinformation spread.

Networked community interactions and ideological homophily

In this study, we identified several networked social influence processes in online misinformation sharing. Online communities are ubiquitous in digital space (Dvir-Gvirsman 2017; Wang, Yang, and Thorson 2021). As individuals interact, basic network-based social processes such as peer influence and social selection occur and influence network outcomes such as misinformation spread (McMillan, Felmlee, and Osgood 2018). In this section we first discusses social processes (transitivity and homophily) that are well documented in the network literature and their implications for understanding misinformation spread. And then, we elaborate on cyberbalkanization, a concept previously studied in other political communication contexts (Holbert, Garrett, and Gleason 2010). We explain how our study contributes to the understanding of cyberbalkanization's influence on misinformation spread.

Our analysis revealed that misinformation spread is more likely to occur among users who share the same political identity. Note that the effect was more stable and robust in the second COVID-19 case. This may be because the dominant majority of users in the first case were identified as conservative, thus reducing the level of variance in the political homophily measure. The principle of homophily suggests that individuals prefer to interact with others who share similar characteristics, beliefs, and behaviors to their own (McPherson, Smith-Lovin, and Cook 2001). Homophily is the first step to bring pairs of individuals closer through their interactions. Through sharing misinformation, homophilous individuals could bond. Prior research on political misinformation spread has shown that preexistent political ideology strongly impacts misinformation sharing (Jerit and Zhao 2020). Our study provides further evidence showing that political homophily is a salient characteristic that drives community members' misinformation-spreading behaviors.

We also found transitivity to be a salient tendency among members who shared misinformation. It is worth noting that transitivity patterns may not be cyclical in online misinformation spread networks, as suggested in the COVID-19 case. This means that when someone retweeted a user through a third party, they may or may not retweet the other user who indirectly retweeted them. However, the significance of transitive triads, in general, suggests that if a mutual third party retweeted two users, the two users are also more likely to retweet each other. This fundamental network-based social process leads to network closure and clustering (McMillan, Felmlee, and Osgood 2018). Our analysis revealed that this is also a critical social mechanism that propels misinformation to spread in online communities.

Cyberbalkanization

The opposite patterns associated with *embedded authority* and *community loyalty* in the process of online misinformation spread revealed in our study contribute to both misinformation research and the literature on cyberbalkanization. Networked social interactions lead to the emergence of particular network structures such as dense clusters at the community level, known as cyberbalkanization (Chan and Fu 2017). In such communities, members take on different roles, and some members may rise to prominence due to their information-sharing activities.

Previous research on cyberbalkanization characterized member positions only using one network structure; however, we measured cyberbalkanization in two different ways that capture users' differing roles in information flows. This is an important point that this study differs from previous research (Chan and Fu 2017). We recognize that some members of cyberbalkanized communities possess high embedded authority and enjoy exclusive popularity within certain communities. In contrast, others target particular communities to share information and demonstrate high levels of community loyalty. While both types of users would be considered similarly in previous research, our study shows that they are different conceptually and exerted different levels of impact in online misinformation spread.

Members with high embedded authority are likely to have strong power within their communities and foster misinformation spread due to networked social influence. The cyber "echo chamber" effect may amplify the impact of embedded authority because members in homophilous networks are more likely to follow the views of popular in-group members (Colleoni, Rozza, and Arvidsson 2014). However, this is not saying that members with high embedded authority do not have strategic intentions. In fact, popular actors in social networks may purposefully influence others after realizing their position in the network (Ibarra 1993). However, members with high embedded authority can impact other community members because of the power brought by their embeddedness in the social network, not because of their intention.

Quite on the contrary, members with high community loyalty concentrate their information sharing activities in certain communities. They may be easily identified as actors practicing strategic manipulations if they have not built their influence in the community via social processes. These members may join a community to bombard the community with certain (mis)information. Although we did not measure their intention, previous literature suggests that such behaviors are often driven by strategic purposes (Chalmers and Shotton 2016). Our analysis suggests that despite their focused misinformation sharing, messages distributed by people with high community loyalty are less likely to be relayed by other community members. One possible explanation is that members in online communities care about their commitment to the group and are reluctant to help strategic outsiders to disseminate their message unless they are also popular among their peers. Together, these results suggest that compared to strategic information manipulations that are not supported by the community's network structure, misinformation shared by members who have high embedded authority in the social network is more likely to gain momentum and fuel wide misinformation spread among the community members.

Overall, our findings provide a nuanced understanding of the mechanism of cyberbalkanization on misinformation and demonstrate the differential effects of networked social influence and strategic information manipulation on online misinformation spread. We show that not all cyberbalkanized individuals contribute to misinformation spread as they may impact misinformation dissemination through different social mechanisms. Depending on the direction of their information exchange activity, some users are significantly more likely to drive the dissemination of misinformation due to the power of their network position, while others are less influential despite their demonstrated devotion to the community.

This finding has implications for policymakers and social networking services (e.g., Facebook, YouTube). Massive amounts of misinformation flood online communities during elections or major crises (e.g., the COVID-19 global pandemic). Our analysis shows that one approach to combat the problem is identifying users who are the authorities within a community and have a history of sharing misinformation. It would be beneficial for social media platforms to censor such accounts or target fact-checking efforts on their messages (tag their messages as potentially misleading to warn other users). Considering there are millions of users in online communities and even the largest technology companies have limited human resources, this approach may be much more targeted and effective than random interventions.

Bots

Our study shows that bots are active participants in the online misinformation dissemination, consisting of approximately 20–30% of the participants in both communities that we studied. Yet, the impact of bots on online misinformation spread is inconsistent and may depend on the context. In the anti-refugee misinformation case, bots significantly promoted the spread of the misinformation. However, in the COVID-19 case, messages of bots were not widely shared by community members. One possible explanation for this divergence is that the first online community is highly homogeneous, with conservative-leaning users as the majority. Yet in the COVID-19 case, the users were more diverse in terms of their manifested political ideologies. It could be that strategic targeting and seeding of misinformation by bots are more effective in homogeneous online communities.

Together, our study demonstrates that networked social influence is consistently more effective than strategic information manipulation to induce high levels of misinformation spread within one interconnected group. Strategic information manipulation may be important in starting the chain of misinformation spread in more homogeneous communities (Vargo, Guo, and Amazeen 2018; Weeks and Gil de Zúñiga 2021). However, they need to be coupled with networked social influence to fuel further misinformation spread (Shin et al. 2018).

Limitations and future directions

This study has several limitations that can be addressed in future research. Although we have identified the different roles between individuals with high embedded authority and high community loyalty, we did not examine how these two network positions interact with each other. Future studies can examine whether members score high on one cyberbalkanization index will be more likely to score high on the other. As communities evolve, it is likely that members' roles are dynamic. Members with high embedded authority and high community loyalty may swap roles over time. For instance, because high community loyalty individuals constantly share information within the community, they may have a higher chance of being perceived as valuable information sources and gradually become embedded authorities. If this occurs, it could indicate that individuals with a strong intention to spread misinformation can become influential by constantly sharing information within a community. Future research could use longitudinal data to explore the dynamics and threshold for such potential role transformation.

Second, we focused on two relatively small core misinformation networks constructed using users' retweets sent in a given period. The sizes of communities and the cutoff dates of the tweets selection window limit the types and lengths of interaction patterns we could observe in this study. By only focusing on two cases of misinformation spread selected, the generalizability of some findings might also be limited. The spread of each piece of misinformation has its unique political and cultural context. Thus, researchers might observe the effect of political homophily and religious identity differently in future cases. For instance, the first anti-refugee case we selected illustrated Islamophobic discourses, and the second COVID-related misinformation case displayed individuals' distrust in science and political authorities. Our control variable, Judeo-Christian identity, showed a significant negative effect in case one but not case two, which indicates that the individuals' prior beliefs have varied effects on perceiving different misinformation cases. Future studies may identify larger networks in different contexts to see if similar interaction patterns exist. Studies can also compare factors that influence members' decisions to share misinformation and scientific information.

Third, our analysis was based on cross-sectional data and only focused on the retweeting relationships. The misinformation spread is a complex process as one piece of misinformation may continue to evolve and develop more layers based on the original message. Future research should track how social influence and misinformation processes play out in online communities over time. Additionally, community-level misinformation spread is a process that is likely influenced by users' socio-psychological characteristics and community interaction dynamics. Social dynamics beyond the observed retweet networks, such as following-follower relationships, remain unknown in our analysis. Future studies may identify members of such communities and use survey or interview methods to explore members' motivations and intentions to spread misinformation. Researchers may also want to study communities formed based on multiple types of relationships and examine whether the two cyberbalkanizition indices (*embedded authority* and *community loyalty)* still predict the flow of the information.

Fourth, there are limitations in our operationalization of the variables. For the political ideology variable, we used users' profile bio to determine their political inclination and coded individuals who did not provide explicit information as nonpartizan. Users who were coded as nonpartizan might identify themselves or indicate their political attitudes in their tweets. Future studies may consider conducting sentiment analysis on users' previous tweets to determine the political inclination of Twitter users. For the bot variable, we detected bots using the algorithm Bot-hunter. However, bot detection on social media remains a challenging issue for practitioners and researchers. What is the best method to detect social media bots is still a matter of debate. We applied one of the newest bot detection algorithms that fit the characteristics of our data and were shown to be highly accurate in detecting Twitter bots. Future studies should explore different bot detection methods and examine bot behavior in different social media contexts. Lastly, bots in the current research are defined and identified as automated accounts, which include malicious bots designed to disseminate misinformation and benevolent bots that automatically relay news (Haustein et al. 2016). Future research should distinguish between various types of bots on social media and examine how they may influence online (mis)information spread differently.

Conclusion

In conclusion, the study reveals that networked social influence has stronger associations with misinformation spread in online communities compared to strategic information manipulation. The nuanced findings regarding the two social mechanisms and of cyberbalkanization – embedded authority and community loyalty – especially warrant further research. The updated understanding of misinformation spreading in online communities could lay an important foundation for developing a holistic approach to combat online misinformation spread. Continued research on this area will help to fully reveal how networked social influence and strategic manipulation interact to shape misinformation-spreading outcomes.

Acknowledgement

We would like to thank Janet Fulk, Peter Monge, Emilio Ferrara, Lindsay Young, members of the Annenberg Networks Network, two anonymous reviewers, and the editor for their constructive feedback on the article.

Funding

The research is supported by the 2016 Dean's Big Data/Social Networks Research Grant from Annenberg School for Communication and Journalism, University of Southern California.

ORCID

Lichen Zhen http://orcid.org/0000-0002-5495-6302
Aimei Yang http://orcid.org/0000-0003-3756-7812

References

Bennett, W. L., and S. Livingston. 2018. The disinformation order: Disruptive communication and the decline of democratic institutions. *European Journal of Communication* 33 (2):122–39. doi: 10.1177/0267323118760317.

Beskow, D. M., and K. M. Carley. 2018. Bot-hunter: A tiered approach to detecting & characterizing automated activity on Twitter. Accessed October 16, 2022. http://www.casos.cs.cmu.edu/publications/papers/LB_5.pdf.

Bessi, A., and E. Ferrara. 2016. Social bots distort the 2016 U.S. Presidential election online discussion. *First Monday* 21 (11). https://firstmonday.org/ojs/index.php/fm/article/view/7090.

Block, P. 2015. Reciprocity, transitivity, and the mysterious three-cycle. *Social Networks* 40:163–73. doi: 10.1016/j.socnet.2014.10.005.

Brainard, L. A. 2009. Cyber-communities. In *International encyclopedia of civil society*, ed. H.K. Anheier and S. Toepler, 587–600. New York: Springer.

Burt, R. S. 2005. *Brokerage and closure: An introduction to social capital*. Oxford, UK: Oxford University Press.

Centers for Disease Control and Prevention. 2020. Reducing stigma. Accessed October 16, 2020. https://www.cdc.gov/coronavirus/2019-ncov/daily-life-coping/reducing-stigma.html.

Chalmers, A. W., and P. A. Shotton. 2016. Changing the face of advocacy? Explaining interest organizations' use of social media strategies. *Political Communication* 33 (3):374–91. doi: 10.1080/10584609.2015.1043477.

Chan, C., and K. Fu. 2017. The relationship between cyberbalkanization and opinion polarization: Time-series analysis on Facebook pages and opinion polls during the Hong Kong Occupy Movement and the associated debate on political reform. *Journal of Computer-Mediated Communication* 22 (5):266–83. doi: 10.1111/jcc4.12192.

Chen, E., K. Lerman, and E. Ferrara. 2020. Tracking social media discourse about the COVID-19 pandemic: Development of a public coronavirus Twitter data set. *JMIR Public Health and Surveillance* 6 (2):e19273. doi: 10.2196/19273.

Cinelli, M., W. Quattrociocchi, A. Galeazzi, C. M. Valensise, E. Brugnoli, A. L. Schmidt, P. Zola, F. Zollo, and A. Scala. 2020. The COVID-19 social media infodemic. *Scientific Reports* 10 (1):16598. doi: 10.1038/s41598-020-73510-5.

Colleoni, E., A. Rozza, and A. Arvidsson. 2014. Echo chamber or public sphere? Predicting political orientation and measuring political homophily in Twitter using big data. *Journal of Communication* 64 (2):317–32. doi: 10.1111/jcom.12084.

Dechêne, A., C. Stahl, J. Hansen, and M. Wänke. 2010. The truth about the truth: A meta-analytic review of the truth effect. *Personality and Social Psychology Review* 14 (2):238–57. doi: 10.1177/1088868309352251.

Del Vicario, M., A. Bessi, F. Zollo, F. Petroni, A. Scala, G. Caldarelli, H. E. Stanley, and W. Quattrociocchi. 2016. The spreading of misinformation online. *Proceedings of the National Academy of Sciences of the United States of America* 113 (3):554–9. doi: 10.1073/pnas.1517441113.

DeMarzo, P. M., D. Vayanos, and J. Zwiebel. 2003. Persuasion bias, social influence, and unidimensional opinions. *The Quarterly Journal of Economics* 118 (3):909–68. doi: 10.1162/00335530360698469.

Dixon, G. N., B. W. McKeever, A. E. Holton, C. Clarke, and G. Eosco. 2015. The power of a picture: Overcoming scientific misinformation by communicating weight-of-evidence information with visual exemplars. *Journal of Communication* 65 (4):639–59. doi: 10.1111/jcom.12159.

Doran, D., H. Alhazmi, and S. S. Gokhale. 2013. Triads, transitivity, and social effects in user interactions on Facebook. In *2013 Fifth International Conference on Computational Aspects of Social Networks*, 68–73. New York: IEEE.

Dvir-Gvirsman, S. 2017. Media audience homophily: Partisan websites, audience identity and polarization processes. *New Media & Society* 19 (7):1072–91. doi: 10.1177/1461444815625945.

Ferrara, E. 2017. Disinformation and social bot operations in the run up to the 2017 French presidential election. *First Monday* 22 (8). https://firstmonday.org/ojs/index.php/fm/article/view/8005.

Freelon, D., and C. Wells. 2020. Disinformation as political communication. *Political Communication* 37 (2):145–56. doi: 10.1080/10584609.2020.1723755.

Guo, L., J. A. Rohde, and H. D. Wu. 2020. Who is responsible for Twitter's echo chamber problem? Evidence from 2016 U.S. election networks. *Information, Communication & Society* 23 (2):234–51. doi: 10.1080/1369118X.2018.1499793.

Guo, L., and C. Vargo. 2020. "Fake news" and emerging online media ecosystem: An integrated intermedia agenda-setting analysis of the 2016 U.S. presidential election. *Communication Research* 47 (2):178–200. doi: 10.1177/0093650218777177.

Hameleers, M., and T. G. L. A. van der Meer. 2020. Misinformation and polarization in a high-choice media environment: How effective are political fact-checkers? *Communication Research* 47 (2):227–50. doi: 10.1177/0093650218819671.

Haustein, S., T. D. Bowman, K. Holmberg, A. Tsou, C. R. Sugimoto, and V. Larivière. 2016. Tweets as impact indicators: Examining the implications of automated "bot" accounts on Twitter. *Journal of the Association for Information Science and Technology* 67 (1):232–8. doi: 10.1002/asi.23456.

Hjorth, F., and R. Adler-Nissen. 2019. Ideological asymmetry in the reach of pro-Russian digital disinformation to United States audiences. *Journal of Communication* 69 (2):168–92. doi: 10.1093/joc/jqz006.

Hochschild, J. L, and K. L. Einstein. 2015. *Do facts matter? Information and misinformation in American politics*. Norman, OK: University of Oklahoma Press.

Holbert, R. L., R. K. Garrett, and L. S. Gleason. 2010. A new era of minimal effects? A response to Bennett and Iyengar. *Journal of Communication* 60 (1):15–34. doi: 10.1111/j.1460-2466.2009.01470.x.

Holland, P. W., and S. Leinhardt. 1971. Transitivity in structural models of small groups. *Comparative Group Studies* 2 (2):107–24. doi: 10.1177/104649647100200201.

Ibarra, H. 1993. Network centrality, power, and innovation involvement: Determinants of technical and administrative roles. *Academy of Management Journal* 36 (3):471–501. doi: 10.2307/256589.

Iribarren, J. L., and E. Moro. 2011. *Affinity Paths* and information diffusion in social networks. *Social Networks* 33 (2):134–42. doi: 10.1016/j.socnet.2010.11.003.

Jang, S. M., T. Geng, J.-Y Q. Li, R. Xia, C.-T. Huang, H. Kim, and J. Tang. 2018. A computational approach for examining the roots and spreading patterns of fake news: Evolution tree analysis. *Computers in Human Behavior* 84:103–13. doi: 10.1016/j.chb.2018.02.032.

Jerit, J., and Y. Zhao. 2020. Political misinformation. *Annual Review of Political Science* 23 (1):77–94. doi: 10.1146/annurev-polisci-050718-032814.

Jost, J. T., S. van der Linden, C. Panagopoulos, and C. D. Hardin. 2018. Ideological asymmetries in conformity, desire for shared reality, and the spread of misinformation. *Current Opinion in Psychology* 23:77–83.

Kasprak, A. 2020. The origins and scientific failings of the COVID-19 "bioweapon" conspiracy theory. Accessed October 16, 2022. https://www.snopes.com/news/2020/04/01/covid-19-bioweapon.

Koch, T., and T. Zerback. 2013. Helpful or harmful? How frequent repetition affects perceived statement credibility. *Journal of Communication* 63 (6):993–1010. doi: 10.1111/jcom.12063.

Krishna, A. 2017. Motivation with misinformation: Conceptualizing lacuna individuals and publics as knowledge-deficient, issue-negative activists. *Journal of Public Relations Research* 29 (4):176–93. doi: 10.1080/1062726X.2017.1363047.

Lai, G., and O. Wong. 2002. The tie effect on information dissemination: The spread of a commercial rumor in Hong Kong. *Social Networks* 24 (1):49–75. doi: 10.1016/S0378-8733(01)00050-8.

Lazer, D. M. J., M. A. Baum, Y. Benkler, A. J. Berinsky, K. M. Greenhill, F. Menczer, M. J. Metzger, B. Nyhan, G. Pennycook, D. Rothschild, M. Schudson, S. A. Sloman, C. R. Sunstein, E. A. Thorson, D. J. Watts, and J. L. Zittrain. 2018. The science of fake news. *Science* 359 (6380):1094–6. doi: 10.1126/science.aao2998.

Lee, T. K., Y. Kim, and K. Coe. 2018. When social media become hostile media: An experimental examination of news sharing, partisanship, and follower count. *Mass Communication and Society* 21 (4):450–72. doi: 10.1080/15205436.2018.1429635.

Lombard, M., J. Snyder-Duch, and C. C. Bracken. 2002. Content analysis in mass communication: Assessment and reporting of intercoder reliability. *Human Communication Research* 28 (4):587–604. doi: 10.1111/j.1468-2958.2002.tb00826.x.

Magelinski, T., D. Beskow, and K. M. Carley. 2019. Graph-Hist: Graph classification from latent feature histograms with application to bot detection. Accessed October 16, 2022. http://arxiv.org/abs/1910.01180

McMillan, C., D. Felmlee, and D. W. Osgood. 2018. Peer influence, friend selection, and gender: How network processes shape adolescent smoking, drinking, and delinquency. *Social Networks* 55:86–96. doi: 10.1016/j.socnet.2018.05.008.

McPherson, M., L. Smith-Lovin, and J. M. Cook. 2001. Birds of a feather: Homophily in social networks. *Annual Review of Sociology* 27 (1):415–44. doi: 10.1146/annurev.soc.27.1.415.

Mejias, U. A., and N. E. Vokuev. 2017. Disinformation and the media: The case of Russia and Ukraine. *Media, Culture & Society* 39 (7):1027–42. doi: 10.1177/0163443716686672.

Milman, O. 2020. A quarter of all tweets about climate change are produced by bots. Accessed October 16, 2022. https://grist.org/climate/a-quarter-of-all-tweets-about-climate-change-are-produced-by-bots/.

Mocanu, D., L. Rossi, Q. Zhang, M. Karsai, and W. Quattrociocchi. 2015. Collective attention in the age of (mis)information. *Computers in Human Behavior* 51:1198–204. doi: 10.1016/j.chb.2015.01.024.

Monge, P. R., and N. S. Contractor. 2003. *Theories of communication networks*. Oxford, UK: Oxford University Press.

Pasek, J., G. Sood, and J. A. Krosnick. 2015. Misinformed about the Affordable Care Act? Leveraging certainty to assess the prevalence of misperceptions. *Journal of Communication* 65 (4):660–73. doi: 10.1111/jcom.12165.

Robins, G., P. Pattison, Y. Kalish, and D. Lusher. 2007. An introduction to exponential random graph (p*) models for social networks. *Social Networks* 29 (2):173–91. doi: 10.1016/j.socnet.2006.08.002.

Roth, Y., and N. Pickles. 2020. Bot or not? The facts about platform manipulation on Twitter. Accessed August 22, 2021. https://blog.twitter.com/en_us/topics/company/2020/bot-or-not.

Shin, J., and K. Thorson. 2017. Partisan selective sharing: The biased diffusion of fact-checking messages on social media. *Journal of Communication* 67 (2):233–55. doi: 10.1111/jcom.12284.

Shin, J., L. Jian, K. Driscoll, and F. Bar. 2017. Political rumoring on Twitter during the 2012 U.S. presidential election: Rumor diffusion and correction. *New Media & Society* 19 (8):1214–35. doi: 10.1177/1461444816634054.

Shin, J., L. Jian, K. Driscoll, and F. Bar. 2018. The diffusion of misinformation on social media: Temporal pattern, message, and source. *Computers in Human Behavior* 83:278–87. doi: 10.1016/j.chb.2018.02.008.

Shulman, S. 2011. DiscoverText: Software training to unlock the power of text. In *Proceedings of the 12th Annual International Digital Government Research Conference: Digital Government Innovation in Challenging Times*, 373. New York: ACM.

Song, H., and H. G. Boomgaarden. 2017. Dynamic spirals put to test: An agent-based model of reinforcing spirals between selective exposure, interpersonal networks, and attitude polarization. *Journal of Communication* 67 (2):256–81. doi: 10.1111/jcom.12288.

Southwell, B. G., and E. A. Thorson. 2015. The prevalence, consequence, and remedy of misinformation in mass media systems. *Journal of Communication* 65 (4):589–95. doi: 10.1111/jcom.12168.

Sunstein, C. R. 2007. Neither Hayek nor Habermas. *Public Choice* 134 (1-2):87–95. doi: 10.1007/s11127-007-9202-9.

Thorson, E. 2016. Belief echoes: The persistent effects of corrected misinformation. *Political Communication* 33 (3):460–80. doi: 10.1080/10584609.2015.1102187.

Vargo, C. J., L. Guo, and M. A. Amazeen. 2018. The agenda-setting power of fake news: A big data analysis of the online media landscape from 2014 to 2016. *New Media & Society* 20 (5):2028–49. doi: 10.1177/1461444817712086.

Vosoughi, S., D. Roy, and S. Aral. 2018. The spread of true and false news online. *Science* 359 (6380):1146–51. doi: 10.1126/science.aap9559.

Vraga, E. K., and L. Bode. 2017. Using expert sources to correct health misinformation in social media. *Science Communication* 39 (5):621–45. doi: 10.1177/1075547017731776.

Waisbord, S. 2018. Why populism is troubling for democratic communication. *Communication, Culture and Critique* 11 (1):21–34. doi: 10.1093/ccc/tcx005.

Wang, L., A. Yang, and K. Thorson. 2021. Serial participants of social media climate discussion as a community of practice: A longitudinal network analysis. *Information, Communication & Society* 24 (7):941–59. doi: 10.1080/1369118X.2019.1668457.

Wang, P., G. Robins, P. Pattison, and E. Lazega. 2013. Exponential random graph models for multilevel networks. *Social Networks* 35 (1):96–115. doi: 10.1016/j.socnet.2013.01.004.

Wang, X., and Y. Song. 2020. Viral misinformation and echo chambers: The diffusion of rumors about genetically modified organisms on social media. *Internet Research* 30 (5):1547–64. doi: 10.1108/INTR-11-2019-0491.

Wardle, C. 2017. Fake news. It's complicated. Accessed October 16, 2022. https://firstdraftnews.org:443/latest/fake-news-complicated/.

Watts, D. J. 2002. A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences of the United States of America* 99 (9):5766–71. doi: 10.1073/pnas.082090499.

Weeks, B. E. 2015. Emotions, partisanship, and misperceptions: How anger and anxietymoderate the effect of partisan bias on susceptibility to political misinformation. *Journal of Communication* 65 (4):699–719. doi: 10.1111/jcom.12164.

Weeks, B. E., and H. Gil de Zúñiga. 2021. What's next? Six observations for the future of political misinformation research. *American Behavioral Scientist* 65 (2):277–89. doi: 10.1177/0002764219878236.

Willson, M. 2010. Technology, networks and communities: An exploration of network and community theory and technosocial forms. *Information, Communication & Society* 13 (5):747–64. doi: 10.1080/13691180903271572.

Wojcik, S., S. Messing, A. Smith, L. Rainie, and P. Hitlin. 2018. Bots in the Twittersphere. Accessed October 16, 2022. https://www.pewresearch.org/internet/2018/04/09/bots-in-the-twittersphere/.

World Health Organization. 2020. How to report misinformation online. Accessed December 27, 2020. https://www.who.int/campaigns/connecting-the-world-to-combat-coronavirus/how-to-report-misinformation-online.