

The Influence of Interdependence in Networked Publics Spheres: How Community-Level Interactions Affect the Evolution of Topics in Online Discourse

Aimei Yang,¹ Ian Myoungsu Choi,² Andrés Abeliuk,³ and Adam Saffer ⁴

¹Annenberg School for Communication & Journalism, University of Southern California, Los Angeles, CA 90089, USA

²Information Sciences Institute, Viterbi School of Engineering, University of Southern California, Marina Del Rey, CA 90292, USA

³Department of Computer Science, University of Chile

⁴Hubbard School of Journalism and Mass Communication, University of Minnesota, Minneapolis, MN 55455, USA

Investigations of networked public spheres often examine the structures of online platforms by studying users' interactions. These works suggest that users' interactions can lead to cyberbalkanization when interlocutors form homophilous communities that typically have few connections to others with opposing ideologies. Yet, rather than assuming communities are isolated, this study examines community-level interactions to reveal how communities in online discourses are more interdependent than previously theorized. Specifically, we examine how such interactions influence the evolution of topics overtime in source and target communities. Our analysis found that (a) the size of a source community (the community that initiates interactions) and a target community (the community that receives interactions), (b) the stability of the source community, and (c) the volume of mentions from a source community to a target community predicts the level of influence one community has on another's discussion topics. We argue this has significant theoretical and practical implications.

Lay Summary

Political discussions online, especially those in the United States, seem to range between harmonious discussions of likeminded people and heated debates that end with few, if any, who have changed their minds. Researchers have often examined these balkanized/polarized situations by studying online communities as isolated echo chambers of opinion. Our study focuses on the interactions between online communities who have different worldviews. We examine communities engaged in the global refugee crisis. We consider how the inter-community interactions influence the agenda of the respective communities. Our longitudinal analysis on the one hand

Corresponding author: Aimei Yang; e-mail: aimei.yang@usc.edu

Editorial Record: First manuscript received on 21 March 2020; Revisions received on 1 November 2020; Accepted by Prof. Sonja Utz on 1 November 2020; Final manuscript received on 2 November 2020

confirms previous studies, namely that intra-community interactions indeed resemble echo chambers. On the other hand, we also find that there is interdependence in the inter-community discussion topics, albeit some communities had greater influence on other communities' discussion topics. For example, larger, more stable communities command more influence.

Keywords: Networked Public Spheres, Cyberbalkanization, Twitter Discourse, Community-Level Interaction, Network Analysis, Dynamic Topic Modeling

doi:10.1093/jcmc/zmab002

Discourse about politics and public affairs is thought to be a fundamental building block of democracy because it can help publics make sense of complex social and political issues, enhance the value of news and information, and encourage civic and political engagement. Historically, in many Western countries the episodes of public discourse about immigration have led to mass mobilizations, affected election outcomes, shaped public policies, and impacted immigrants' opportunities and challenges in their host countries (Waisanen, 2012). The same appears to be true for the ongoing international refugee crisis that has forcibly displaced 70.8 million people at last count (UNHCR, 2019). Yet, in contrast to previous episodes, much of the current discourse about refugees and immigration rights is taking place in networked public spheres.

The scholarly attention on networked public spheres has tended to focus on how the communication environments create structures that influence individuals and their discussions (Benkler et al., 2015). Some scholars have lauded online platforms for providing networked infrastructures that facilitate political discourse and actions (Marwick & Boyd, 2011). Most of such studies emphasize networks as conduits for accessing and disseminating information and as the space for various types of actors to participate in "open" discussions (Ausserhofer & Maireder, 2013). Meanwhile, online platforms can create network structures that contribute to cyberbalkanization and the fragmentation of networked public spheres. For instance, when "people seek out only like-minded others and thereby close themselves off from ideological opposition, alternative understandings, and uncomfortable discussions," as Brainard (2009, p. 598) notes, they create fragmented network structures. Evidence indicates fragmented networked public spheres occur from selective exposure and repeated interactions among like-minded others (Benkler et al., 2015). While cyberbalkanization and polarized opinions within communities are well-researched (Freelon et al., 2018), little is known about whether interactions among communities can influence their discussion topics.

Recognizing that the opinion environment within a community tends to be homophilous, this study examines how interactions among communities online, especially communities with opposing views, affect each other's discussion topics. We draw from research on networked public spheres (Benkler et al., 2015), cyberbalkanization (Brainard, 2009), and the evolution of topics (Ausserhofer & Maireder, 2013; Griffith & Mullins, 1972; Uitermark et al., 2016) to hypothesize that although the interactions among communities do little to change people's view, they may influence the agenda for what other communities talk about. To test our hypotheses, our computational research design combines social network analysis, topic modeling, community detection, and temporal topic evolution models. The findings show how community level interactions connect the fragmented public sphere into a mega network where siloed communities still influence other communities' discourse topics.

Fragmented networked public spheres

The public sphere—the space between the state and society—is “a constellation of communicative spaces in society that permit the circulation of information, ideas, debates—ideally in an unfettered manner—and also the formation of political will (i.e., public opinion)” (Dahlgren, 2009, p. 34). The original Habermasian notion of a public sphere captures the interactions between members of a society and how such interactions sustain democracies (Habermas, 1962). Yet, even before the digital era, the public sphere was unlikely to be a singular, homogenous arena for discourse deprived of hierarchies. Fraser (1990) has highlighted the significant exclusion of publics based on gender, social classes, and ethnicity. As the dominant class, often heterosexual men, assume universal class in the public sphere, their hegemonic public sphere could block others’ access, and force counterpublics to create alternative public spheres (Dahlgren, 2009).

The proliferation of social and digital media and related fragmentation have increased the dispersion of public discourse and offered “citizens most invisible in mainstream politics radical new potentials for identity negotiation, visibility, and influence” (Jackson & Foucault Welles, 2016, p. 399). Rather than a singular public sphere, later work, including that of Habermas’ (2015) own reflections, suggest the multiplicity of public arenas, and the embracing of multidimensional, heterogenous spheres (Benkler et al., 2015; Dahlgren, 2009).

Contemporary works have conceptualized *networked public spheres* as “the range of practices, organizations, and technologies that have emerged from networked communication as an alternative arena for public discourse, political debate, and mobilization alongside, and in interaction with, traditional media” (Benkler et al., 2015, p. 596). Such understanding recognizes that individuals can choose whom they meet and talk with online. As such, they tend to cluster according to their identities, ideologies, and other preferences, and form communities that offer sociability, support, and a sense of identity (Himmelboim et al., 2013; Jackson & Foucault Welles, 2016; Marwick & Boyd, 2011). Such interaction patterns could lead to homogeneous groups that essentially create “echo chambers” within communities (Colleoni et al., 2014). Studies have confirmed that the networked public spheres are characterized by clusters of densely connected social-mediated communities that are linked by sparse ties (Brainard, 2009). Social-mediated communities can provide the space where individuals and organizations engage in collaborative actions, information seeking or sharing, and creating shared meaning, values, and identities (Marwick & Boyd, 2011).

Cyberbalkanization and the formation of online communities can mutually reinforce each other (Freelon et al., 2018). Cyberbalkanization refers to the phenomenon in which users interact primarily with like-minded others and avoid ideological opposition and uncomfortable discussions (Brainard, 2009). Cyberbalkanization tends to be especially salient when information exchange is related to politicized issues such as the global refugee crisis (Chan & Fu, 2018). Cyberbalkanization facilitates the formation of online communities through providing members with shared meaning, identities, and opportunities for repeated interactions. This means community members form frequent interactions within their communities while avoiding outsiders with disagreeable views. On the other hand, as clusters of communities contribute to a fragmented network structure, the silos of opinion environments could reinforce cyberbalkanization, especially within groups with extreme worldviews. Cyberbalkanization and the fragmentation of public spheres can mutually-reinforce members’ worldviews.

For example, Himmelboim et al. (2013) examined networks among Twitter users who discussed President Obama’s 2012 State of the Union speech. The study found that users tended to cluster into communities and were more likely to follow others from their own cluster than from other clusters.

Further analysis of users' discourse revealed that users tune in to a narrow segment of the wider range of politically oriented information sources. These results align with [Brummette et al.'s \(2018\)](#) research. They studied the network structure of connected users who discussed "fake news" and found users tend to form ideologically homophilous clusters, use "fake news" to disparage the opposition, and condemn information disseminated by the opposing party.

Previous research on the structure of discourses in networked public spheres confirms that users tend to form homophilous communities. Frequent interactions within these communities appear to reinforce existing worldviews ([Himmelboim et al., 2013](#)). Yet, we recognize that communities are not necessarily simply isolated from one another. Taking a more macro view, one can see that different communities—even those communities that are ideologically opposed—are interconnected as part of a networked sphere of public discourse. Our question here, then, is what is the missing link that connects the fragmented networked public sphere?

While the current study focuses on social mediated communication, scholars in other fields have explored similar questions such as how interactions among scientific communities influence each other's research topics and contribute to scientific evolutions ([Griffith & Mullins, 1972](#); [Uitermark, et al., 2016](#)). We draw from these areas to explore whether ideologically distinct communities can set the agenda for what other communities talk about.

Contentious community interactions and topic evolution

To understand the mechanisms and potential impact of community level contentious interactions on the evolution of public discourse, we draw upon [Griffith and Mullins's \(1972\)](#) work on contentious groups in scientific communities and how they affect scientific revolution. In scientific fields, revolutionary changes are often advanced by a small group of scholars who collaborate closely with like-minded colleagues to challenge dominant paradigms or rival groups. [Griffith and Mullins \(1972\)](#) described a process that contributes to the rise of coherent groups that leads to revolution.

Specifically, in scientific communities, coherent groups tend to have contentious differences in their perspectives compared to other groups, and therefore draw boundaries between groups. Once a boundary is established, some groups place themselves at the center of a debate by attracting negative attention from more established counter-groups. Second, strong support or positive ties within groups contribute to the group's solidarity and provide a support system that enables further revolutionary works. Third, these coherent groups tend to have iconic leaders who perform the critical intellectual leadership roles to represent the group to its members and outsiders. The leaders effectively muster more support to advance their perspectives. Generally, group level interactions influenced other groups' research agenda.

[Uitermark et al. \(2016\)](#) applied [Griffith and Mullins's \(1972\)](#) insights in their study of contentious public discourses in the Netherlands about the integration of ethnic minorities and proposed several expectations about the community-level interactions between opposing movements. First, they suggested that the public interested in a particular issue would cluster into communities, with some communities confronting others. Importantly, they understood the contentious and coherent communities as a network mechanism that could facilitate interactions among like-minded publics and cultivate shared identities and perspectives on an issue. [Uitermark et al.](#) also suggested that communities need to draw negative responses from a countermovement in order to catapult them into the center of a debate and gain agenda setting power. Here it is assumed that eliciting negative criticism alone does not sustain influence; instead, communities need to win support within their own

communities or they may be dismissed as pariahs and face exclusion. Ultimately the cohesion and stability of a community is significant because it provides “a cohesive base of followers whose supportive relations countervail criticisms” (Uitermark et al., 2016, p. 108).

Informed by previous research (Griffith & Mullins, 1972; Uitermark et al., 2016), we propose that both the *characteristics* of interacting communities and their *interaction patterns* affect their level of influence on other communities’ discussion topic evolutions. Unlike previous studies, we consider the communication content as a dynamic rather than static texts. The semantic content of public discourse can be analyzed through topic based community detection such as Latent Dirichlet Allocation (LDA) and its extensions (Blei & Lafferty, 2006). Although traditional static topic modeling approaches tend to assume topics are shared interchangeably by all documents, such assumption has been scrutinized. The combination of machine learning and natural language processing allow researchers to find patterns of words in text-based data using hierarchical probabilistic models. The dynamic topic modeling approach has examined the evolution of discourse topics in sequentially organized corpus of text-based data (Blei & Lafferty, 2006).

Characteristics of interacting communities

Two basic community characteristics that we consider here is community size and community stability over time. Griffith and Mullins’s (1972) research found that newcomer research groups needed to grow to a certain size in order to commend certain levels of influence on other groups’ research agendas. Similarly, other studies have found that in the networked public sphere, communities with more members tend to command more influence on the temporal evolution of topics due to their size (Freelon et al., 2018). Communities with more members can amplify voices through their collective follower or friend bases and disproportionately influence topic evolution. Therefore, we propose H1:

H1: In networked Twitter spheres, the sizes of both source (H1a) and target (H1b) communities will positively associate with their level of influence on temporal topic changes.

Griffith and Mullins (1972) discussed the importance of coherent and loyal supporters. Research on the Advocacy Coalition Framework (Sabatier & Weible, 2007) also argues that coalitions of political actors is critical to ensure their influence on public and policy agenda. Because coherent coalitions appear stronger in public debate, sub-communities of coalitions often coordinate their arguments in support of each other. In the context of social mediated communication, the effect may be even more salient. For communities to influence discourse topics over a period of time, if these communities themselves are not very stable, it might be difficult to imagine they can achieve any lasting impact on discussion topics. As such, we propose the following hypothesis:

H2: In networked Twitter spheres, community stability for both source (H2a) and target (H2b) communities will positively associate with their level of influence on temporal topic changes.

Community interaction patterns

In terms of community interaction patterns, we propose that the valence of interactions and the frequency of interaction (one community mentions the other) may matter. A key insight from Griffith and Mullins (1972) and Uitermark et al.’s (2016) research is that negative interactions among communities could drive public attention and discourse. In other words, the communities on opposing sides of the refugee issue that engage in negative interaction could exert greater impact on the temporal changes of discourse topics. Another insight from Uitermark et al.’s (2016) research is that positive interactions among like-minded communities is crucial for boosting the visibility of these

communities and therefore increases their influence on temporal topic changes. Based on this rationale, we propose that the following hypothesis:

H3: In networked Twitter spheres, the valence of interactions from the source (H3a) to target (H3b) communities will significantly influence temporal topic changes in these communities.

Additionally, the frequency of interactions among communities may also drive topic evolution. Research on issue attention cycle (Downs, 1972) and “political waves” suggest that at any given time, public spheres could only accommodate a limited number of issues (topics) and actors need to compete for attention. The more interactions occur among two communities, the more likely that they can set the agenda for each other’s discussion topics. Therefore, we propose:

H4: In networked Twitter spheres, source (H4a) and target (H4b) communities that engage in frequent interactions will significantly influence temporal topic changes in these communities.

Method

Data

Twitter has more than 330 million users worldwide. The United States is the country with most active users (68 million users, accounting for 22% of the U.S. population) worldwide. Studies have shown that Twitter provides a salient platform for public discourse of important social and political issues (Ausserhofer & Maireder, 2013; Benkler et al., 2015). Based on these reasons, we select Twitter to conduct our research. Using DiscoverText (Shulman, 2011),¹ we gathered tweets mentioning various forms of “refugee” or “illegal immigrants” in 2016. The data collection took two steps. First, we queried LexisNexis and gathered the number of articles published with the above keywords each day in 2016. Then we identified 26 news spikes, which are periods when the amount of media coverage mentioning refugees was two standard deviations above average. Assuming Twitter discussions could influence the media coverage and vice versa, we gathered tweets one day before a news spike and gathered tweets a day after to capture tweets that proceeded or followed the media coverage, respectively. In total, data from 86 days were collected from the year of 2016.

We preprocessed the data set with the following steps. First, non-English tweets were filtered out. Next, to exclude non-active users, the tweets whose user’s status count² was equal to or lower than 10 were removed. After filtering, the data set contains 723,267 users who generated 2,264,341 tweets mentioning “refugee” or “illegal immigrants.” We further extracted retweeting relationships to form a retweet network. A retweet network is a directed graph in which a node is each user and an edge represents a retweet relation.³ The retweet network consists of 546,989 users and is connected by 1,143,590 edges.

Analytic procedures

Community detection

Next, the Louvain algorithm⁴ was applied for community detection. This algorithm was performed on the retweet network using eight different resolution parameters for searching the optimal modularity.⁵ As a result, 318 communities were detected with 0.701 of modularity value. The top 10 communities consist of 367,709 users, which accounted for 72% of the total number of users (510,506) in all the communities detected. And the users from top 10 communities contributed 1,502,361 tweets,

which accounted for 82.9% of the total number of tweets (1,812,457) from all the communities detected (See [Appendix B](#) for details about the structures of the 10 communities). Therefore, we chose to conduct further study mainly on the top 10 communities. In addition, the PageRank⁶ values of all the members within each community were computed in order to identify the influencers in each community. Further, each community was named with the top influencer's Twitter ID.

Valence-based interactions among communities

To identify valence-based community-level network interactions, the sentiment analysis method was applied. We used VADER⁷ because it was developed specifically for social media contexts using human-validated sentiment lexicon. VADER was developed by applying the following procedures. First, an extensive sentiment lexicon list was built by incorporating the existing well-established sentiment lexicon list (LIWC, ANEW, and GI) and additional lexical features such as emoticons, acronyms, and initialisms. Second, intensities of lexicons from the list were rated on a scale from -4 to +4 (extremely negative to extremely positive) by human raters. Lastly, generalizable heuristics, which were based on grammatical and syntactical perspectives (e.g., exclamation, capitalization, contrastive conjunction, etc.) were identified and incorporated into the method.

The compound score which we used to identify the valence of a tweet in this study is calculated by summing the valence scores of each sentiment lexicon in a tweet and then normalized to be between -1 and 1 (extremely negative to extremely positive). The compound score is categorized into positive if the score is $\geq .05$ and negative if the score is $\leq -.05$, and otherwise neutral. We applied sentiment analysis as follows. First, sentiment scores of all tweets were computed. Second, the tweets mentioning users in different communities were grouped by the source community and the target community.⁸ Lastly, the average sentiment scores of all 10 communities were computed.

Document influence model

We used the document influence model⁹ (Gerrish & Blei, 2010) to quantify the influence of each community in discourse topic evolution. This model allows us to classify the documents into corresponding topic distributions and it allows us to capture the influence of the past document on the future documents in terms of a certain topic T , which is represented as l_T . Applying this model involved several steps as follows.

As a preprocessing step, we performed a tweet pooling scheme, under which all the users in the same community can be considered as co-authors. First, all the retweets were excluded for the redundancy reason and transformed into lowercase. Second, mention tags (@userid), URLs, tokens that consist of only non-alphanumeric characters, short tokens (Length < 3 characters), short tweets (The number of tokens < 3) were removed. Next, lemmatization was performed with tweets from the above steps. During the lemmatization step, punctuations and stop words¹⁰ were removed and preserved only as a noun, adjective, verb, adverb. Finally, tweets from the above steps were grouped by the same community and concatenated into one document of each community.

After the preprocessing step, we trained the LDA¹¹ model with three different numbers of topics (5, 10, 15) on 10 concatenated documents from the top 10 communities based on the number of members in each community in order to find the most relevant topics among the top 10 communities. As a result, $k=5$ was chosen as the optimal number of topics which results in the most relevant and consistent topics. Lastly, the document influence model was trained. In this step, all the tweets from one community were concatenated as the documents on a daily basis. As a result, 172 documents from each community, therefore, 1,720 documents from the top 10 communities were obtained and trained with $k=5$ as the topic number parameter.¹²

Dependent variable

Our dependent variable is pairwise influence score. From the document influence model (Gerrish & Blei, 2010), we computed the influences of the documents in terms of topic evolution and also topic distributions of the documents. We introduced the pairwise influence score $I_{c_1c_2}$, which could show us how much community c_1 has an influence on community c_2 in terms of discourse topic evolution. $I_{c_1c_2}$ is computed by summing up the multiplication of l_{c_1T} and a_{c_2T} over all the topic T . l_{c_1T} is the average influence score of all the documents from the community 1 as to the evolution of topic T over time. a_{c_2T} is the average topic distribution of topic T from all the documents of the community 2 over time, which can be interpreted as the attention of community 2 as to topic T over time.

$$I_{c_1c_2} = \sum_{T \in \text{topics}} l_{c_1T} a_{c_2T} \quad (1)$$

Independent variables¹³

Sentiment scores

Sentiment scores from the source community to the target community were calculated by taking the average VADER sentiment scores of all the tweets in which the *source community mentioned the target community* over a given time unit. Sentiment score from the target to the source community is similar to the previous variables but calculates the average VADER sentiment scores of all the tweets in which the *target community mentioned the source community* over a given time unit.

Normalized mention counts

Next, mentioned counts from the source community to the target community take the mentioned counts of all the tweets in which the source community mentioned the target community over a given time unit. Additionally, the mentioned counts from above were divided by the total mentioning volume of the source community over a given time unit in order to remove the effect of the size of the source community. Again, the next variable is similar but distinct because it calculates mentioned counts from the target community to the source community by summing all the tweets in which the target community mentioned the source community was calculated over the given time unit. Likewise, the mentioned counts from the target community to the source community were normalized in the same way as above.

Community stability score

Following Metzler, Günemann and Miettinen, (2019), we computed the containment index for each community to assess communities' stability over time. The containment index is calculated using the following equation:

$$C(G, H) = \frac{G \cup H}{\min(G, H)}$$

Where G and H stands for two sets of nodes at two time points. This index indicates which percentage of nodes from the smaller set is contained in both sets. The higher the overlap, the more stable a community is. Community membership of the users was computed by the Louvain algorithm, which

was applied only once over the whole year. Therefore, it assumes that the membership of users does not change during the year. Next, we look at whether a user was active in a time window. The containment index was calculated from two member sets of a community (G and H) from consecutive time units.

Size of communities

This variable considers the number of users in the (a) source communities, and (b) target communities. The number of users in the community are counted over a given time unit.

Results

Descriptive statistics about top communities and topics

Our exploratory analysis identified the top communities, top topics, and the relationship between each top community and top topics. Most communities are dominated by the presence of a handful of top influencers.¹⁴ In some cases such as the AJ+ (@ajplus) community, this Qatari state-funded media conglomerate alone dominated the community. The analysis of the document influence model revealed the most influential topics. [Table 1](#) reports the top five most influential topics, names and description, topic distribution scores, and examples of keywords. The most influential topic is topic 1: refugee rights. This is also the topic that has the largest proportion of the topic distribution. This topic accounted for more than 65% in all other top 10 communities' topic distribution. Tweets on this topic focused on how countries can help with refugee and promoted awareness campaign for World Refugee Day, and popular hashtags in this topic also expressed positive sentiment toward refugees such as #withrefugee and #refugeewelcome.

The second most influential topic is topic 3: U.S. election fraud and refugee terrorists. Tweets under this topic focused on how refugees coming through the Mexico–U.S. border are terrorists, and specifically linked Hillary Clinton with illegal immigrants to suggest election fraud. The third most influential topic is topic 2: Syrian refugee resettlement, which discusses resettlement issues and challenges facing millions of Syrian refugees, especially children, who currently live in countries such as Turkey, Jordan, and Lebanon. The fourth most influential topic is topic 4: UN sexual abuse of refugees, which alleges that staffers in UN led refugee resettlement programs sexually harassed female and children refugees. Topic 5 the last most influential topic: Hungary and Germany border dispute. This topic focused on Germany's criticism on Hungary, which ultimately led to Germany shutting down and implementing "border controls" to refugees.

Additionally, [Table 2](#) provides average influence scores of the top 10 communities and some basic information about the most influential accounts in each community. Among the top 10 influencers, there are media personalities and official accounts such as Fox Nation hosts (@DiamondandSilk), Qatar media (@ajplus), Russian media (@RT_com), and two UK media (@guardian and @mailonline). There are leaders from the NGO community such as UNICEF (@UNICEF), UNHCR (@UNrefugees), and Kenneth Roth (@KenRoth, head of Human Rights Watch). Politicians like @SushmaSwaraj and celebrities like @JohnFugelsang who advocated for refugee rights were also found to be influential. The influences score suggests that the size of the communities might be one of the factors affecting the influence scores since the order of the rank of the influence score is similar to that of the size of the community.

In order to examine how each community contributed to the popularity of the top five topics, we conducted topic distribution analysis. The topic distribution proportions can be interpreted as the amount of tweets from each community dedicated towards certain topics. [Appendix I](#) shows how

Table 1 Five Topics and Examples of Keywords in the Document Influence Model

Topic	Topic Names & Description	Topic Probabilities	Examples of Keywords
Topic 1	Refugee rights: Discussed how countries can help refugee; refugee rights are human rights	0.748	Worldrefugeeday, Australia, auspol, withrefugee, refugeeewelcom, Australian, nauru, human, right, thank
Topic 2	US election fraud and refugee terrorists: Misinformation about how refugees coming through the U.S.–Mexico border are terrorists who will help Hillary Clinton to win election	0.102	Obama, Trump, border, American, Hillary, resettlement, Mexico, terrorist, immigration
Topic 3	Syrian refugee resettlement: Discussed the resettlement and challenges associated with millions of Syrian refugees currently living in Turkey, Jordan, and Lebanon	0.097	Turkey, Syrian, Greece, child, Lebanon, Jordan_violation, rights_amman, Greek_island nato, coverage, Macedonia
Topic 4	UN sexually abused refugee: Discussed how UN peacekeepers sexually abused refugees. The information was provided by the Heritage Foundation, a conservative organization supporting Trump.*	0.030	Reporting_girl, media_defend, heritage_foundation, science_taxe, raise_trillion, program_stalled, surge_sex
Topic 5	Hungary vs. Germany border dispute: Discussed how Hungary closed border to refugees while Germany slams disgraceful EU	0.022	Germany, Merkel, German, border, Palestinian, Denmark, Swedish, flow, influx, Berlin, shelter, Hungarian, Israel

Note: Topics are manually labeled for convenience. *The information is available here: <https://www.heritage.org/civil-society/commentary/the-un-sex-scandal>

much influence each community has on each topic. For the two topics that are relatively unfriendly towards refugees or organizations advocating for refugees (topic 3: U.S. Election and refugee as terrorists & topic 4: UN sexually abused refugee), the @DimondandSilk community and @Mailonline community contributed most to their popularity. The NGO communities (UNICEF, UNHCR and Human Rights Watch) were mostly contributing to topic 1: refugee rights. The Russian media led

Table 2 Average Influence Scores of Top 10 Communities in 2016

Rank	Community	Description of Account	Community Size	Influence Score
1	DiamondandSilk	Two vloggers who claimed to be President Donald J Trump's most loyal supporters. They are also hosts of Fox Nation.	76,704	0.00395
2	UNICEF	The United Nations Children's Fund, was created by the United Nations General Assembly to provide emergency food and healthcare to children and mothers in countries devastated by war.	57,191	0.00234
3	Guardian	<i>The Guardian</i> is a British daily newspaper. <i>The Guardian</i> is part of the Guardian Media Group, owned by the Scott Trust.	55,796	0.00231
4	UNrefugees	The United Nations High Commissioner for Refugees is a United Nations program with the mandate to protect refugees, forcibly displaced communities and stateless people, and assists in their voluntary repatriation, local integration or resettlement to a third country.	42,674	0.00177
5	KenRoth	Kenneth Roth, an American attorney, has been the executive director of Human Rights Watch since 1993.	32,120	0.00168
6	RT_com	RT is a Russian international television network funded by the Russian government.	31,172	0.0013
7	MailOnline	MailOnline is the website of the <i>Daily Mail</i> , a newspaper in the United Kingdom, and of its sister paper <i>The Mail on Sunday</i> .	22,225	0.00123
8	JohnFugelsang	John Joseph Fugelsang is an American actor, television personality and comedian.	18,078	0.00099
9	Ajplus	AJ+ is an online news and current events channel run by Al Jazeera Media Network, which is a Qatari state-funded global media conglomerate.	16,132	0.00091
10	SushmaSwaraj	Sushma Swaraj was an Indian politician and a Supreme Court lawyer.	15,617	0.00086

Note: Means are zero normalized and listed descending order by influence score.

Table 3 Multiple Regressions Results With Pairwise Influence Score as the Dependent Variable

Variable		β
Intercept		.001
<i>Community Characteristics</i>		
Community Size	Size of the source community (H1a)	.42***
	Size of the target community (H1b)	-.07**
Community Stability Score	Community stability of the source community (H2a)	.27***
	Community stability of the target community (H2b)	.13***
<i>Community Interaction Patterns</i>		
Sentiment Score	Sentiment score from the source community to the target community	-.03
	Sentiment score from the target community to the source community	-.04
Mention Counts	Mentioned counts from the source community to the target community	-.01
	Mentioned counts from the target community to the source community	.07**
R^2		.41
Adjusted R^2		.40

Note. All β are standardized coefficients for multiple regression. * $p < .05$. ** $p < .01$. *** $p < .001$.

community showed great interest in the topic 5: Hungary vs. Germany border dispute and drove the discussion on topic 2: Syrian refugee resettlement.

Hypotheses testing

To test our hypotheses, we examined how community size and stability, valence of interaction and volume of interaction (number of mentions) affect the temporal topic evolution. Table 3 provides detailed results from the model.¹⁵ Together, these interactions explain 40% of variance in temporal topic changes. This provided consistent support for H1_{a-b} across many conditions. Specifically, the sizes of the source and target community were statistically significant ($\beta = .421^{***}$, $\beta = -.07^{**}$) on the pairwise influence score. Put another way, when a large sized community discusses certain topics and has members who directed parts of the conversation at another community whose size is small, the topics of the source community could set the discussion agenda in the target community.

H2_{a-b} directs attention to the effect of community stability. For communities to achieve a lasting impact on what other communities are talking about, they need to remain stable (a significant portion of members stayed in the same community over one year) over time. This is true for the source community. When a source community wants to influence the discussion topic in a target community, it needs to stay relatively stable. However, in turn, when a target community wants to achieve a lasting impact on what members of the source community discuss, they do not need to stay relatively stable and coherent over time. H2_a was supported and H2_b rejected.

H3_{a-b} considered that the valence of interactions from the source and target communities might influence topic changes. Neither positive nor negative interactions between a source and a target community appear to have made consistent difference in affecting topic changes. Therefore, the results provided no support for H3_{a-b}.

H4_{a-b} are concerned with whether the volume of interactions among source (a) and target (b) communities influenced the other communities' discussion topics. While the volume of interactions from the source community does not seem to affect topic changes, under most conditions the volume of the interactions¹⁶ from target to source communities yield significant results and therefore provide partial support for H4_b.

Overall, the analysis showed that for one community to influence the discourse topics in another community, the community would have a better chance if it is a large community and if it stays stable over time. Additionally, when a target community engages back, it might also shape what the source community discusses.

Discussion

Previous cyberbalkanization studies have focused on how mediated connections are built among homogenous actors and how their opinions tend to polarize within communities (Freelon et al., 2018; Himelboim et al., 2013). In this study, we shifted attention away from *within* communities and instead focused on the interdependence among communities. We explored how community level interactions affect temporal changes in communities' discussion topics. This focus allows us to consider how topics from one community set the agenda in others. Our analysis shows the interdependence among communities is pervasive and impactful. Communities can influence discussion topics of other communities when they meet several key conditions.

Revealing a fragmented networked public sphere

Research on networked public sphere has mapped the structure of public discourses on social media platforms (Brummette et al., 2018; Jackson & Foucault Welles, 2016). These studies have documented a fragmented public sphere, where like-minded actors cluster into tight communities. Our study found similar results. Our analysis identified the top 10 communities with relatively stable structure and dense connections within each community. This finding revealed the social structure that enabled cyberbalkanization (Freelon et al., 2018). It is through the repeated, clustered interactions, selective exposure can be sustained over time (Benkler et al., 2015).

In addition, our analysis revealed that over time, five dominant topics captivated the public attention (accounted for more than 65% of public discourse). While all communities discussed these topics at some point, their attitudes towards these topics are fairly different. Our analysis showed that conservative communities tend to be hostile to refugee-friendly causes while progressive communities and NGO centered communities tend to be friendly to refugees. This finding is consistent with previous cyberbalkanization research (Brainard, 2009) and predictions derived from the Advocacy Coalition Framework (Sabatier & Weible, 2007). The explanation here is that actors with different ideologies and political agendas functioned liked advocacy coalitions in the networked public sphere. These coalitions promoted messages that were consistent with their worldview and use coordinated efforts in the process of advocacy.

We also found a strong elite dominance within each community. The top influencers in the top communities were INGOs, IGOs, politicians, celebrities, and media (see [Table 2](#)). These individuals/organizations are iconic in the sense that they attracted the lion share of attention and won support from their community members. Indeed, such finding is consistent with expectations derived from [Griffith and Mullins's \(1972\)](#) work. Griffith and Mullins theorized that in contentious contexts, elites are necessary to ensure internal solidarity within communities while coordinating attacks on opposing movements. Additionally, these elites rather than everyday people emerge as top influencers may also reflect the nature of the research context as an international affair. In terms of the ongoing refugee crisis, elites often have better access to firsthand information. This finding highlights that when studying different issues, it is critical to identify “localized influencers” that are specific to each community. In contentious social issues, online communities may provide a localized context for influencers to exert their impact. As such, community detection may be the necessary first step for influencer identification and theorization. Future studies may also explore the role of localized influencers in keeping communities stable. Research can examine if the removal of these localized influencers would significantly reduce the stability of their communities.

So far, our study has confirmed the existence of communities on Twitter in which users were tightly connected. These communities tend to have iconic leaders who are exceptionally influential when compared to other members. These communities were able to coordinate their discussion efforts in a coalition-style over a one-year period to attack topics they disapproved and promoted topics they favored. More importantly, as we will further discuss below, these communities were not isolated from each other. Indeed, 18% of the links are between users of different communities (See [Appendix B](#) for the relative densities of the communities as compared to the overall density in the network). Communities did interact with each other, though not always with a positive sentiment. The interactions among communities is the link missing in previous research that ties the fragmented networked public sphere into a connected sphere.

The missing link in the fragmented networked public sphere

Our study showed that when communities engage with each other, topics that one community cared about could potentially set the discussion agenda in other communities, depending on a few conditions. This does not mean that community level interactions could change different community members' attitudes towards issues, but does suggest such interactions changed the focus of public attention (what people talked about). Research on issue attention cycle ([Downs, 1972](#)) suggest that issues need to compete for public attention in the crowded public sphere. Our study suggests that if members of a community want to engage others on issues they care about, they should interact more with members of other communities, even communities they completely disagree with. Such interactions appear to be effective at setting the discussion topics agendas of other communities and potentially driving the discourse topics of the overall public sphere. This finding offers implication for understanding of the “echo chamber” effect ([Colleoni et al., 2014](#)). We found that although “echo chambers” exist within communities, when considering the overall discourse on refugees, there are conversations between the political left and right. This finding is consistent with [Barberá et al.'s \(2015\)](#) study, which analyzed over 150 million tweets and found that conversations on political issues at the national level are more dynamic than simply polarized around ideological clusters.

Additionally, we found that several conditions are critical for achieving and sustaining inter-community influence. The first condition is community size. Our results indicate that the size of the source communities had a profound impact on the discourse topic evolution in the target community.

The effect would be especially good if the size of the target community is relatively small. Taken together, large size communities that frequently initiate attacks or approvals at another community may enjoy considerable influence on discourse topics in the other community. As expected, communities of large size tend to be more successful when pushing their agenda. In our case study, partly due to the size of the @DiamondandSilk community and its keen interests in xenophobia rhetoric, one of the most influential topics in the data set are anti-refugees.

The second condition is community stability. Our analysis showed that some of the largest communities in our data set are also quite stable. Not only did members consistently tweet about refugee issues over a one year period, they also interacted consistently with roughly the same group of people. One possible explanation is that these community members are issue publics who cared a great deal about the refugee issue. In addition, cyberbalkanization may have provided them a sense of community (Freelon *et al.*, 2018). Our analysis showed that this stability is crucial for communities to sustain their influence on topic evolution. The finding is consistent with Uitermark *et al.*'s (2016) prediction that coherent groups of supporters are critical for countermovements to advocate their argument. For advocacy groups, our finding suggests that they need to engage supporters and sustain a stable level of participation.

The third condition is the volume of interactions. Griffith and Mullins's seminal work and its later development (Uitermark *et al.*, 2016) conceptualize contentious public discourse from a relational approach, and suggest that the back-and-forth interactions between opposing communities could drive topic changes. Interestingly, our finding suggests that when a target community engages back with a source community, it actually gains the power to influence the source community's discussion topics.¹⁷

The recommendation based on our finding is that some target communities should engage back with consistent messages. Such engagement may not be able to change other communities' viewpoints, but would spread the impact of their topics and expose them to more people. The same recommendation can be offered to strategic efforts aiming at debunking misinformation. For example, communities that wish to share correct information about a global pandemic should actively engage with communities that spread misinformation. Their actions may not necessarily change attitudes but could increase the counterpublics' exposure to facts and set their discourse agenda.

Taken together, our study sheds light on how networks and discourse are interrelated in the networked public sphere. Our study confirms that within communities, cyberbalkanization could hold communities together and lead members to be "closed in" in their opinion environment. Meanwhile, our study also shows that among communities, inter-community interactions connect communities through shared attention to the same topics. Moreover, the meso-level community interactions allow one community to set the agenda of another's discussion topics and may eventually shape public discourse.

Our study offers a nuanced, multilevel perspective on networked public spheres. This perspective recognizes the local clusters and isolated opinion environment individual users may experience while underscores the meso-level connections that tie different communities into an interconnected, interdependent public sphere. This perspective challenges scholars to explore new theories that accommodate the different tensions at the micro and meso levels and opens up new frontiers for research.

Methodological implications, limitations, and future directions

This study offers several methodological contributions. First, we see communication structure and content as inherently intertwined and mutually dependent. In order to simultaneously study the

structure and content of communication, we combined social network analysis, topic modeling, community detection and temporal topic evolution analysis. This method considers communication networks and discourse topics as dynamic processes, and examines how network structure affects topic evolutions. This approach precisely shows how much influence each community had on each topic, and how communities are interconnected through their interactions. This method can be used to study a range of communication topics such as the spreading of misinformation, movement mobilization, and various roles of online communities.

Moreover, the community-level interaction approach also illuminated new research possibilities. Previous research tends to focus on how network positions of individual actors or communication strategies affect social media advocacy. The current study suggests that community structure and inter-community interaction patterns may reveal new mechanisms on how social influence occurs among communities. The study suggests that future social-mediated communication research needs to adopt multilevel and multimodal research perspectives that simultaneously accommodates individual interactions, community interactions, and macro-network level interactions. To isolate analysis within one level may lead to incomplete conclusions or lose sight of critical mechanisms that are only visible through a multilevel design. Moreover, in this study we only examined Twitter. As such, we know little about how other types of media like newspapers influence the interactions among communities. Future studies could examine communities' links to media and how such links impact interactions among communities. Additionally, in this study we quantified community level influence through influence scores. Future research may further validate such influence with additional indicators.

Finally, we also recognize several epistemological and ethical challenges inherent in big data and computational methods. First, when large quantities of data are aggregated and processed, the richness of individual cases may be lost in the process of aggregation. While our study revealed macro and meso level patterns, future studies could benefit from methods such as netnography to document the nuanced daily interactions and cultural routines within and between communities. Second, social media data sets, however big, could only account for specific platforms' user populations. This may raise the issue of representativeness. As these methods may disproportionately highlight and reinforce digital inequality. In the case of current study, while a portion of refugees may have access to the Internet, millions who are in the direst conditions may not. These are issues that need to be counted for in future research.

Supporting information

The following supporting information is available for this article:

Appendix A The Number of Detected Communities and Modularity Values for Each Resolution Parameter.

Appendix B1 Community Densities and Ratios as Compared to the Overall Density in the Network.

Appendix B2 The Cumulative Ratio of the Community Sizes in the Network.

Appendix C The Distribution of the User-Level PageRank in Each Community.

Appendix D Sentiment Interaction Network of the Top 10 Communities.

Appendix E Log Perplexity Values From Three Different LDA Models Based on a Number of Topics.

Appendix F Descriptive Statistics of Independent Variables.

Appendix G Pairwise Pearson Correlation Coefficients of Independent Variables.

Appendix H The Top 10 Influencers in the Top 10 Communities.

Appendix I Topic Distributions of Each of the Top 10 Communities.

Appendix J Results of Multiple Regressions With Pairwise Influence Score as the Dependent Variable. Each Model is Different in Terms of the Number of the Topics Learned in the Document Influence Model (T: The Number of the Topic).

Appendix K Results of Multiple Regressions With Pairwise Influence Score as the Dependent Variable. Each Model is Different in Terms of the Resolution Limit Parameter in the Community Detection Method (R: Resolution Limit Parameter).

Appendix L Results of Multiple Regressions With Pairwise Influence Score as the Dependent Variable. Each Model is Different in Terms of the Number of the Topics Learned in the Document Influence Model (T: The Number of the Topics). Additionally, the Condition in Which the First Community (DiamondandSilk) Was Excluded Was Applied.

Endnotes

1. At the time of data collection, through a paid subscription, DiscoverText offers access to the Premium Twitter API, which offers the ability to collect relevant data directly from the full Twitter firehose.
2. User's status count refers to the number of tweets (including retweets) issued by the user. The removal of inactive users (this step excluded 1.1% users) helps to reduce noise and improve the accuracy of document influence models.
3. The source node and the target node are respectively tweet sender and users who retweet this message.
4. The Louvain algorithm is the most commonly used method for community detection although there can be bias towards low-degree users. The algorithm is based on the process of maximizing the modularity thus leading to a small number of the large size of communities and a large number of the small size of communities. This study used the Networkx package in Python language with a resolution parameter of 1.1.
5. Modularity is a measure indicating how much a certain network or graph has dense connections between the nodes within the same community and sparse connections between the nodes from the different communities. Please see [Appendix A](#) with the results.
6. PageRank is a link analysis algorithm that can show the relative importance of each node in the graph. The node with high PageRank value in the graph means that it is linked by many nodes with high PageRank value. In this study, Networkx package in Python language was used and the used alpha parameter was 0.85. Please see [Appendix C](#) for distribution of user-level PageRank in communities.
7. VADER (Valence Aware Dictionary and sEntiment Reasoner) is a sentiment analysis method developed specifically for social media contexts. This method is based on a human-validated sentiment lexicon. Using this method, the sentiment score in the text can be computed, which leads to positive, neutral, or negative sentiment depending on the score (Positive: score ≥ 0.05 , Negative: score ≤ -0.05 , Neutral: else).
8. Source community: A community in which the users mentioned the other users in different communities in their tweets. Target community: A community in which the users were mentioned by the other users in different communities. [Appendix D](#) visualizes the sentiment interaction network of the top 10 communities.

9. Document influence model: Document influence model is an extension of the dynamic topic model (Blei et al., 2006), which can capture the influence of the past documents towards future documents in terms of topic change.
10. Stop words: Stop words are the most common words, such as, the, is, at, and so on. In this study, NLTK package in Python was used.
11. The LDA model is a generative statistical model that can classify the documents in similar topics. Please see [Appendix E](#) for Log perplexity values from three different LDA models.
12. Gensim package in Python language was used. A used time slice parameter was based on a weekly basis.
13. Please see [Appendix F](#) and [G](#) for descriptive statistics and pairwise Pearson coefficients.
14. See [Appendix H](#) for the top 10 influencers in each of the top 10 communities.
15. In addition, [Appendix J](#), [K](#), and [L](#) provides detailed results from the assessment of the robustness of the model in terms of the number of the topics learned in the document influence model, the resolution parameter settings in the community detection method, and whether or not excluding the first community in the model.
16. The only exception is when the number of topics is set to be larger than five (see [Appendix I](#)). But the effect changes to be positively significant when we drop the largest community (see [Appendix K](#)).
17. The only exception is when we include the largest community, @DiamondandSilk, in the analysis. One possible explanation is that this community is very large and no other community could effectively sets its agenda. It could also be because this community is led by extremely conservative influencers and their discussion topics are resilient to change.

References

- Ausserhofer, J., & Maireder, A. (2013). National politics on Twitter: Structures and topics of a networked public sphere. *Information, Communication & Society*, 16(3), 291–314. doi: 10.1080/1369118X.2012.756050
- Barberá, P., Jost, J. T., Nagler, J., Tucker, J. A., & Bonneau, R. (2015). Tweeting from left to right: Is online political communication more than an echo chamber? *Psychological Science*, 26(10), 1531–1542. doi:10.1177/0956797615594620
- Benkler, Y., Roberts, H., Faris, R., Solow-Niederman, A., & Etling, B. (2015). Social mobilization and the networked public sphere: Mapping the SOPA-PIPA debate. *Political Communication*, 32(4), 594–624. doi:10.1080/10584609.2014.986349
- Blei, D. M., & Lafferty, J. D. (2006, June). Dynamic topic models. In *Proceedings of the 23rd international conference on machine learning* (pp. 113–120). New York: Association for Computing Machinery.
- Brainard, L.A. (2009). Cyber-communities. In H. K. Anheier and S. Toepler (Eds.), *International encyclopedia of civil society* (pp. 587–600). New York: Springer Science & Business Media.
- Brummette, J., DiStaso, M., Vafeiadis, M., & Messner, M. (2018). Read all about it: The politicization of “fake news” on Twitter. *Journalism & Mass Communication Quarterly*, 95(2), 497–517. doi: 10.1177/1077699018769906
- Chan, C. H., & Fu, K. W. (2018). The “mutual ignoring” mechanism of cyberbalkanization: Triangulating observational data analysis and agent-based modeling. *Journal of Information Technology & Politics*, 15(4), 378–387.

- Colleoni, E., Rozza, A., & Arvidsson, A. (2014). Echo chamber or public sphere? Predicting political orientation and measuring political homophily in Twitter using big data. *Journal of Communication*, 64(2), 317–332.
- Dahlgren, P. (2009). *Media and political engagement: Citizens, communication, and democracy*. Cambridge, MA: Cambridge University Press.
- Downs, A. (1972). Up and down with ecology: The “issue-attention” cycle. *The Public Interest*, 28(1), 38–50.
- Fraser, N. (1990). Rethinking the public sphere: A contribution to the critique of actually existing democracy. *Social Text*, 25/26, 56–80.
- Freelon, D., McIlwain, C., & Clark, M. (2018). Quantifying the power and consequences of social media protest. *New Media & Society*, 20(3), 990–1011.
- Gerrish, S., & Blei, D. M. (2010, June). A language-based approach to measuring scholarly impact. *ICML*, 10, 375–382.
- Griffith, B. C., & Mullins, N. C. (1972). Coherent social groups in scientific change. *Science*, 177(4053), 959–964. https://www.jstor.org/stable/1734125?casa_token=qAOwtDTj4BQAAAAA%3AwYdu5CHbcjKCfT8MGsxNJewVY6rNKy3ZLlPLwzcTAOWhNgdlikVe6VpPx15E3QZi9fxYe4YriHY2Co73V9-I7l_fks_FslEEvM9PHc-AplFull9ps4Y&seq=1#metadata_info_tab_contents
- Habermas, J. (1962). *The structural transformation of the public sphere: An inquiry into a category of bourgeois society* (trans. T. Burger). Cambridge, MA: MIT press.
- Habermas, J. (2015). *Between facts and norms: Contributions to a discourse theory of law and democracy*. Maiden, MA: John Wiley & Sons.
- Himmelboim, I., Smith, M., & Shneiderman, B. (2013). Tweeting apart: Applying network analysis to detect selective exposure clusters in Twitter. *Communication Methods and Measures*, 7(3–4), 195–223. doi:10.1080/19312458.2013.813922
- Jackson, S. J. ?0026; & Foucault Welles, B. (2016). #Ferguson is everywhere: Initiators in emerging counterpublic networks. *Information, Communication & Society*, 19(3), 397–418. doi:10.1080/1369118X.2015.1106571
- Marwick, A. E., & Boyd, d. (2011). I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. *New Media & Society*, 13(1), 114–133. doi:10.1177/1461444810365313
- Metzler, S., Gunnemann, S., & Miettinen, P. (2019). Stability and dynamics of communities on online question–answer sites. *Social Networks*, 58, 50–58.
- Sabatier, P. A., & Weible, C. M. (2007). The advocacy coalition framework. In P. A. Sabatier (Ed.), *Theories of the policy process* (pp. 189–220). Boulder, CO: Westview Press.
- Shulman, S. (2011, June). 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* (pp. 373–373).
- Uitermark, J., Traag, V. A., & Bruggeman, J. (2016). Dissecting discursive contention: A relational analysis of the Dutch debate on minority integration, 1990–2006. *Social Networks*, 47, 107–115. doi:10.1016/j.socnet.2016.05.006
- UNHCR. (2019). Figures at a glance: Statistical Yearbooks. Available on <https://www.unhcr.org/en-us/figures-at-a-glance.html>.
- Waisanen, D. (2012). Bordering populism in immigration activism: Outlaw–civic discourse in a (counter) public. *Communication Monographs*, 79(2), 232–255. doi:10.1080/03637751.2012.673230