Social Network Analysis and Mining (2024) 14:236
https://doi.org/10.1007/s13278-024-01384-w

CASE REPORT

Norm propagation in online communities: structural, temporal, and community analysis

Raed Alharbi[1] · Youval Kashuv[2] · Yan Wang[3] · Wenwen Dou[4] · My T. Thai[3]

Received: 19 September 2024 / Revised: 9 November 2024 / Accepted: 12 November 2024 / Published online: 14 December 2024

Abstract

Understanding descriptive positive and negative norms, such as their propagation speed, is crucial for shaping individuals' behaviors towards public health guidelines and designing successful promotion campaigns to encourage positive norms. Unfortunately, conducting in-depth analyses related to community characteristics and the diffusion of descriptive norms is complex and context-dependent, influenced by variables across time such as what the norms are and the structure of the community it belongs to. To address this gap, this paper presents a comprehensive analysis of norm propagation in online communities, such as macro and micro analyses, structural and temporal analyses, and activity-based norm life cycle analyses. We also investigate the community's propagation networks to understand the overlapping influencers. Through these analyses, we aim to reveal the dynamics of norm diffusion, understand influence patterns within communities, and identify influential users and clusters contributing to norm adoption and propagation. Our finding shows that negative norms display a shorter life cycle in contrast to positive norms. Additionally, positive norms demonstrate a longer life cycle, while negative norms display a comparatively shorter duration. We also find that engagement near the norm's disappearance is less frequent in negative norms compared to positive norms, where engagement persists.

Keywords Norms · Social network · Propagation network · COVID-19

1 Introduction

The widespread adoption of social media platforms has significantly influenced the spread of various social norms, both positive and negative. Positive descriptive norms, which define appropriate behavior, serve as guiding principles in society (Chen and Hong 2015). By reinforcing acceptable behavior through mechanisms such as praise or reward, these norms establish standard codes of conduct (McDonald et al. 2014).

Conversely, negative descriptive norms capture behaviors deemed inappropriate or undesirable (Chen and Hong 2015). Unlike positive descriptive norms, norms of this type are reinforced by negative outcomes, for example criticism (Chen and Hong 2015; Hassell and Wyler 2019). The impact of such norms can be viewed through the lens of the COVID-19 pandemic.

Positive descriptive norms advocated for wearing masks, promoting social distancing, and sharing accurate information. Negative norms became apparent through the spread of conspiracy theories, misinformation, and disinformation via social media platforms.

✉ Raed Alharbi
ri.alharbi@seu.edu.sa

Youval Kashuv
youval.kashuv@ufl.edu

Yan Wang
yanw@ufl.edu

Wenwen Dou
wdou1@uncc.edu

My T. Thai
mythai@ufl.edu

1 College of Computing and Informatics, Saudi Electronic University, Riyadh 11673, Saudi Arabia

2 Computer and Information Science and Engineering Department, University of Florida, Gainesville, FL 32611, USA

3 Department of Urban and Regional Planning, University of Florida, Gainesville, FL 32611, USA

4 College of Computing and Informatics, University of North Carolina at Charlotte, Charlotte, NC 28223, USA

Exploring the impact social media had on the emergence and spread of descriptive norms during the COVID-19 pandemic is crucial for a number of reasons. To begin, these norms have a major influence on individuals' responses towards public health guidelines (Neville et al. 2021). As such, understanding how positive descriptive norms can be utilized to improve efforts aimed to mitigate the spread of viruses, such as COVID-19, is essential. Simultaneously, investigating dynamics of negative descriptive norms, provides opportunities to effectively counter the spread of misinformation, mitigating its harmful effects on public health initiatives.

Moreover, analysing norm propagation in macro-level and micro-level networks reveals how descriptive norms spread within online communities (Vosoughi et al. 2018; Shu et al. 2020). As depicted in Fig. 1, these networks are structured hierarchically, where different levels of the network provide unique insights into the propagation of descriptive norms. Macro-level networks track the emergence and diffusion of norms, highlighting the impact of influential users, while micro-level networks illuminate properties concerning norm adoption within local communities.

However, despite the usefulness of such investigations, conducting in-depth analyses of community characteristics and their relation to the diffusion of descriptive norms is a complicated and context-dependent process. This is because norms are subject to many degrees of freedom across time, such as what the norm is and the properties of the global and local communities the norm belongs to. Additionally, the majority of efforts have been dedicated to understanding norms in physical (Cao et al. 2021; Eskytė et al. 2020) or simulated (Hu and Leung 2017) environments. This limits the applicability of such findings to online social networks.

Besides this, little research has been conducted into understanding the driving forces behind dissemination of positive and negative norms across various network levels. The impact of influencers on the dynamics of such norms is another unanswered question. This paper aims to address these gaps in knowledge by examining the following questions:

- **Q1**: Does the propagation of positive and negative norms within hierarchical networks follow distinct patterns throughout their life cycles? If so, what are the temporal dynamics and characteristics of these norms?
- **Q2**: Do influencers in communities that support positive and negative norms exhibit similar patterns? How do these influencers differ in overlapping communities?

Findings: Our investigation reveals the following:

- Negative norms display a shorter life cycle, while positive norms demonstrate longer life cycle.
- Positive norms demonstrate a consistent and slower propagation of tweets, whereas negative norms experience faster dissemination, especially during the early stages of emergence.
- Engagement near the norm's disappearance is less frequent in negative norms. This is in contrast to positive norms, where engagement persists.

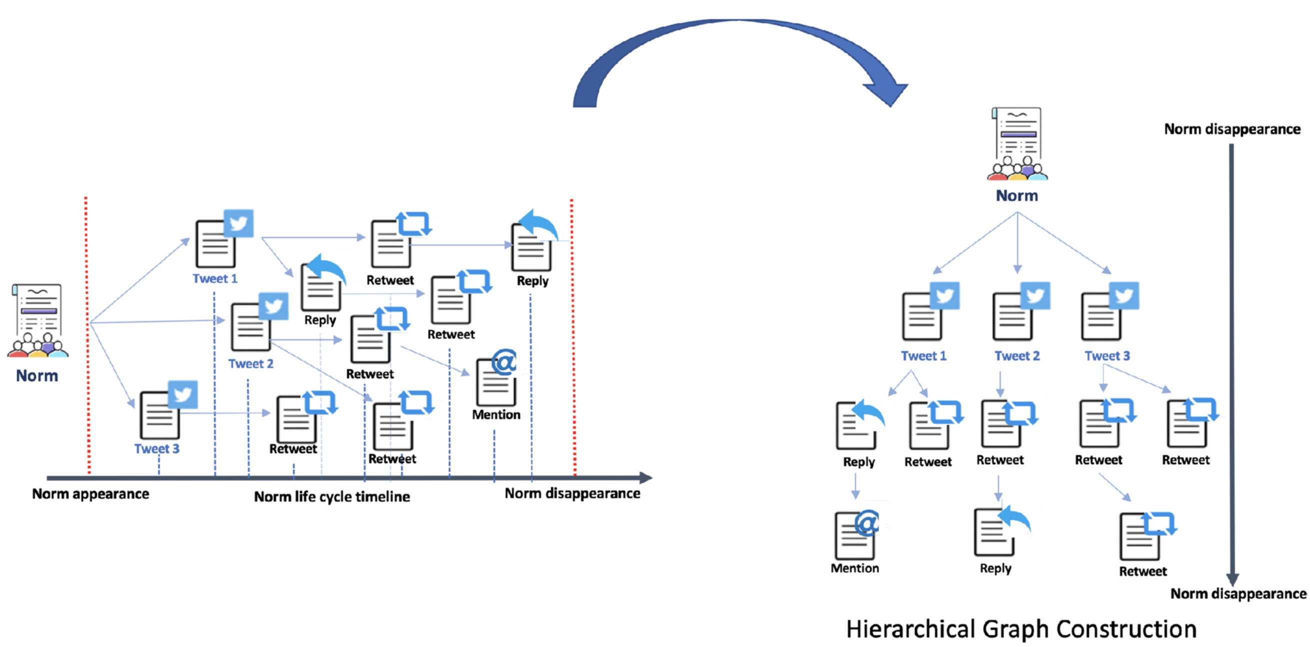

Fig. 1 A visualization of a norm's hierarchical propagation network on a social media platform, from its emergence to its disappearance

- Communities associated with positive norms show higher levels of centralization when compared to those associated with negative norms.

Organization: The remainder of the paper is structured as follows: Sect. 2 presents the related works. Our hierarchical graph construction is then introduced in Sect. 3 and the corresponding in-depth analysis is presented in Sect. 4. Finally, Sect. 5 presents the paper's limitations and conclusion.

2 Related work

2.1 Norm propagation analysis

Recently, the investigation of social norm dynamics within communities has received increased attention, partly because of the COVID-19 pandemic. For example, Latkin et al. (2022) use logistic regression to find associations between social norms and three COVID-19 prevention behaviors: social distancing, washing hands, and wearing face coverings. Using these findings, they suggest that campaigns geared towards public health communication focus on more than simply raising risk perceptions. Sikali (2020) investigates social distancing norms, showing that they increase levels of social rejection, negatively influencing individuals' desire to socialize. Saint and Moscovitch (2021) use a clinical setting to examine the impact of mask-wearing on social anxiety. They hypothesize that the driving forces behind this increase are: (a) individuals' perceptions of social norms related to mask-wearing; (b) individuals' experiences in which masks limit accurate interpretation of social and emotional cues. When it comes to negative norms, Romer and Jamieson (2020) demonstrated that conspiracy beliefs impeded attempts at mitigating the spread of COVID-19. Overcoming these obstacles requires continuous communication efforts from public health officials through means such as traditional media outlets, which have also been prone to support COVID-related conspiracy theories.

Despite extensive efforts towards researching social norm dynamics, several knowledge gaps and challenges remain. First, because most studies focus only on the physical world, our understanding of norm's life cycles in online communities is not yet clear. Second, no studies have explored the temporal aspect of descriptive norms within diffusion-based models, models focused on tracking the propagation of norms throughout social networks. Third, little research has been conducted to understand how influencers impact the dissemination of positive and negative norms. Additionally, the influence of community properties on the propagation of norms in real world settings also remains unexplored. Addressing these knowledge gaps is the first step towards a comprehensive understanding of social norm dynamics within online communities, allowing for the development of effective norm intervention and management strategies.

2.2 Community analysis of propagation networks

Many studies have explored the role of social norms in shaping individuals' beliefs and behaviors in online communities (Savarimuthu and Cranefield 2011; Hawkins et al. 2019). One such case (Hawkins et al. 2019) focuses on norms that provide stable behavioral expectations on the individual and population level. On the individual level, pre-existing norms are examined using advanced cognitive procedures (ex. social reasoning), where as the population level focuses on the structure of social networks and their respective impact on norms.

However, these studies often rely on artificially constructed communities. Since real world communities tend to be much more complex than artificially constructed counterparts, with higher variation in both influencers and community structure, the applicability of such findings to the real world becomes limited. Thus, these gaps must be addressed to provide valuable and usable insights for policymakers seeking to promote positive behavior and combat negative behavior. A thorough investigation of the temporal aspect of norm dynamics, such as the speed of propagation, is furthermore necessary to optimize the timing and effectiveness of norm interventions.

3 Hierarchical graph construction

In this section, we introduce the construction of hierarchical propagation networks for negative and positive norms. In order to track changes in these networks over time and extract insights, we must first define the norm's life cycle. In this study, we formulate the life cycle of a norm to consist of several stages, as depicted in Fig. 1. The initial stage, emergence, represents the first observation. Subsequent tweets expressing agreement with the norm are classified as supporters. Retweets subsequent to the supporters' tweets are then termed distributors, as they are responsible for spreading the views of supporters to a broader audience.

Since the data pre-processing conducted is significant for building the graph used in this study, we first describe it in Sect. 3.1. Section 3.2 then discusses hierarchical propagation networks of positive and negative norms.

3.1 Dataset preparation

To obtain an initial dataset concerning COVID-19 vaccines, we employ the X platform streaming API. The data collection period, starting January 1, 2021 and ending September 30, 2021, focuses on tweets explicitly addressing COVID-19 vaccinations. To achieve this, we use a list of relevant keywords that include "vaccination,""sputnik,"and"vaccine,"a longside specific names of COVID-19 vaccine manufacturers like"Pfizer"and"Moderna,"to filter for relevant tweets. Through this process, we extract 2378 tweets and 9410 retweets containing the keywords about the negative and positive norms from a total of 138, 578..

Identifying online communities that support diverse norms often relies on communication networks that leverage the retweet functionality as a key indicator. This is because retweeting behavior primarily occurs among users who have established follower relationships, serving as a reflection of active social connections between them. Notably, empirical studies examining retweeting behavior in the context of COVID-19 information have revealed that more than half of such retweets originate from users with established follower/following relationships (Gao et al. 2021, 2023). This finding presents an opportunity to measure the social influence carried by individuals on social media, as those who retweet or respond to messages from others are likely influenced by the content they engage with. Consequently, analyzing user interactions such as retweets and replies provides insights into the propagation dynamics of specific norms among social media users, effectively capturing the unique patterns of norm diffusion. Thus, we examine retweeting behaviors to gain an understanding of the social dynamics and influence of online communities.

During the COVID-19 pandemic, various new positive descriptive norms emerged, including social distancing (Allen IV et al. 2021) and wearing masks (Dillard et al. 2021). At the time, these norms were largely accepted by individuals globally, as reflected by the propagation of such positive behaviors on social media platforms. However, the pandemic and social media platforms similarly gave rise to a plethora of negative norms. These included conspiracy theories (Himelboim et al. 2023) such as (i) Bill Gates being involved in the creation of COVID-19 to implant microchips in people, (ii) COVID-19 is a genetically modified organism (GMO), (iii) COVID-19 is a bio-weapon, and (iv) COVID-19 vaccines are untested and poisonous. To select positive of negative norms relevant to our study, we use a formula of"word A + word B". Examples of this are shown in Tables 1 and 2.

3.2 Propagation networks

To properly investigate the dissemination patterns of positive and negative norms, we develop a hierarchical propagation network. This network operates across various levels of granularity, tracing the propagation process of norms from their emergence to disappearance via retweet chains. The network employs directed edges, connecting distributors to supporters, with the edge weight indicating the frequency with which distributors share behaviors expressed by supporters. This formulation effectively represents the spread of the norm in question.

Through this construction, we can efficiently identify influential users and communities involved in the dissemination process. Mathematically, the propagation network can be represented as a directed graph $G = (V, E)$, consisting of a set of nodes representing users V, and a set of directed edges representing retweet relationships E. The weight of an edge w_{uv} reflects the number of times user u retweeted content from user v within our dataset.

Table 1 Positive norms

Topics	Keywords
Maintaining social distance during Pandemic	"distance", "keep distance", "distancing", "6 feet", "6-foot rule", "social distancing", "physical distancing"
Wearing masks after the spread of COVID-19	"Wear a mask", "keep mask on", "WearMask", "Mask On", "wearing mask", "FaceMaskMaskUP", "mandatory masks", "mandatory face mask"

Table 2 Negative Norms

Topics	Keywords
Bill Gates was involved in the creation COVID-19 with the goal of microchipping people	"Bill Gates","microchipping+Bill Gates", "Gates+pandemic simulations"
COVID-19 is genetically modified organism (GMO)	"GMO","genetically modified organism","big pharma","Gates pharma","Fauci pharma","genetically modified"
COVID-19 is a biological weapon	"COVID19+weapon","weapon covid=19","biological weapon"
COVID-19 vaccines are untested and poisonous	"haven't been tested","not tested","skip+trail", "poison","isn't be tested","wasn't tested", "aren't tested","didn't be tested","doesn't be tested"

4 Norms analysis

The objective of this section is to utilize various analytical techniques to examine different aspects of norm dynamics. The first part, structural analysis, focuses on identifying the underlying structural aspects of norm emergence and dissemination. In the second part, the macro and micro analysis explores the spread of norms in hierarchical networks, considering the macro-level patterns of norm diffusion among communities and the micro-level complexity of interactions between individual users.

The following part, temporal analysis, investigates social norms dynamics from a temporal viewpoint, differentiating between how positive and negative norms change over time. In the last part, we explore activity-based aspects of norms, focusing on the life cycle of norms through the lens of user actions and behaviors. This approach allows us to further distinguish between the dynamics positive and negative norms.

4.1 Structural analysis

Within the constructed propagation network, we can effectively capture the dissemination patterns of both positive and negative norms within our dataset, providing insights into the individuals responsible for sharing these patterns. We begin by examining the structural properties of the network to better understand the pattern in which a norm spreads on a global level. Within our structural analysis, we examine a variety of graph properties, with a heavy focus on the distribution of node degrees. This enables a better understanding of the role influencers play during different periods of a norm's life cycle. Overall, structural properties of the constructed propagation network provide insight into what factors influence the dissemination of norms.

Figure 2 depicts the propagation network of positive and negative norms, with the left graph representing positive norms and the right graph representing negative norms. Within these graphs, darker shades of nodes correspond to higher degrees, while darker edges indicate higher edge weight. The lightest edges correspond to sharing behavior that occurred four times. By examining these two graphs, it becomes evident that negative norms display a higher presence of influencers when compared to the propagation network of positive norms. This indicates that the propagation of negative norms is heavily reliant on influencers. Additionally, the strong connections in the positive norms graph primarily occur between influencer and non-influencer nodes. This is not the case in the negative norms graph, where influencers tend to establish strong connections with other influencers in their community. This discrepancy may indicate that positive norms rely on authoritative sources (Gadzekpo et al. 2023), such as the CDC (Varma et al. 2023), in order to spread. In contrast, negative norms, such as conspiracy theories, tend to be reinforced only within tight-knit communities already biased towards them (Dow et al. 2021).

To better understand the implications of network structure on positive and negative norms, we turn our attention to analyzing the degree distribution of influential nodes. Denoted $P_{deg}(k)$, the value of this discrete distribution is the fraction of nodes in the graph with degree k, and it has a strong influence on the rate at which a norm spreads. Our analysis reveals that nodes in positive norm networks have a slightly higher degree than their counterparts, indicating that networks fostering positive norms tend to be more centralized.

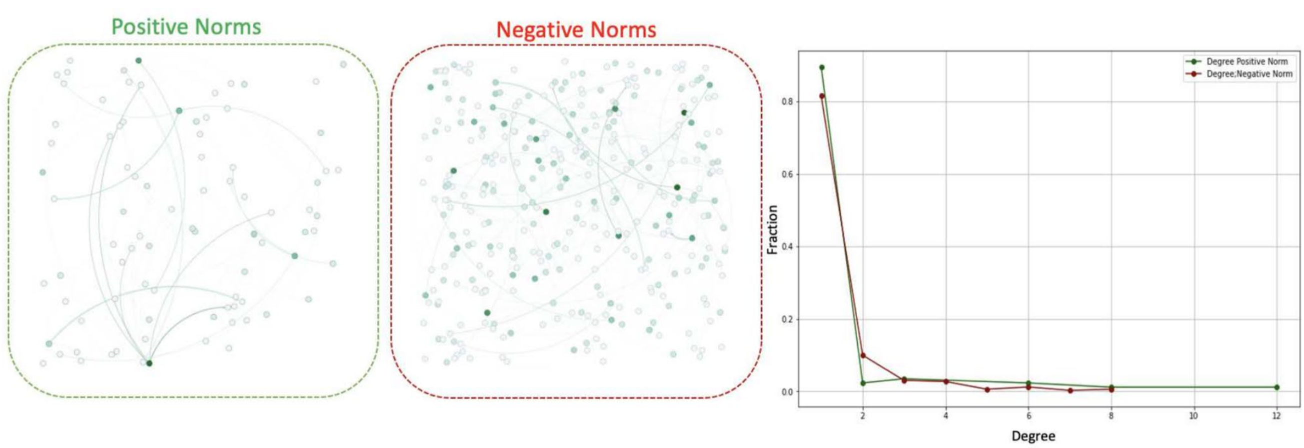

Fig. 2 The propagation network of positive and negative norms, along with the corresponding degree distribution. In the graph associated with positive norms, strong connections primarily connect influencers to non-influencer nodes. Conversely, in the negative norms graph, connections are more frequent, and strong connections can be observed between influencers and other influencer nodes

4.2 Temporal analysis

To provide an exhaustive understanding of norm propagation dynamics, the temporal aspect of user interactions, such as the frequency and intensity of user retweets over time, must also be taken into account. Such analyses may be implicitly captured when modelling temporal patterns using neural networks (Wang et al. 2020), but what these black-box models gain in complexity, they lose in explainability. For these analyses to be useful to policy makers, the underlying rational behind learned features must be made clear.

To address this limitation, we extract several explicit features concerning the temporal aspect of norm dynamics. These features are carefully crafted to underscore distinct properties of norm propagation and investigate whether variations exist between positive and negative norms. Specifically, we extract the following time-dependent features:

- **C1:** Average time elapsed since a norm's appearance and the subsequent tweets supporting it. This feature sheds light on the short-term support that follows a norm's appearance and may provide important indications of users' immediate reactions.
- **C2:** Average time elapsed since instances of norm support and the following distribution of supporting tweets. By measuring the average time a norm takes to spread, we can gauge its rate of propagation through the network. This offers valuable insights, quantifying the speed of propagation and the corresponding norm's intensity.
- **C3:** Interval between consecutive distributors, indicating the time gap between successive actions of distributors during norm propagation. This feature offers key information surrounding the intensity in which distributors are actively engaging in propagating the norm, highlighting the speed of information flow throughout the network.
- **C4:** Time difference between a norm first supporting post and the action of the last distributor. This feature captures the norm's life cycle, measuring the elapsed time since its initial support to its eventual disappearance. Analyzing this duration provides insights into the overarching dynamics of norm propagation, including its duration and longevity.

The temporal attributes of positive and negative norm graphs, depicted in Fig. 3, highlight key differences between positive and negative norms. To begin, looking at **C1** and **C2**, it can be observed that although both positive and negative norms exhibit quick initial support, distributors react significantly quicker for negative norms. This may indicate that for negative norms such as conspiracy theories to spread and survive, they must be exposed to as most people as possible before fact checking can take place. Looking at **C3**, we can observe that distributors engage for longer periods of time in the case of negative norms. This finding could also represent the flooding of misinformation in an attempt to overshadow the truth. It may also indicate the use of bots to target specific users.

Turning our attention to **C4**, it is revealed that negative norms display a shorter life cycle in comparison to positive norms, meaning that on average, negative norms persist for shorter durations. Interpreted together with **C1**, **C2**, and **C3**, these observations could further indicate that negative norms must be spread quickly in order to survive, and counteracting misinformation and disinformation could be as easy as flooding the network with the truth.

These results underscore key differences between the temporal aspects of positive and negative norm dynamics. The shorter life cycle, rapid support and lengthened distribution of negative norms separate their dissemination from that of positive norms, and provide insight into taking action against them.

4.3 Activity-based norm life cycle analysis

To provide a comprehensive understanding of norms and their life cycles, their activity-based aspects must also be taken into account. Activity-based analyzes provide insights into the varying rates at which different norms spread, capturing short-term dynamics and enabling the detection of irregularities within certain time frames. We adopt uniform time intervals, as illustrated in Fig. 4, to provide a straightforward format for comparative assessment.

To further differentiate between the dynamics of positive and negative norms, we dedicate this section to tracking and analyzing the sharing behaviors of norms across different time frames. Specifically, we divide a norm's life cycle into the following three periods:

- Period 1, emergence, captures the norm's initial establishment phase.
- The second period, spanning three months, representing a norm's stability.
- The third and final period, disappearance, emphasizes a norm's gradual decline until elimination.

We evaluate the impact of different types of norms by tracking how often norm-related posts are shared during specific time periods. This approach allows us to understand how norms evolve over time and measure their influence within each phase. In select, in particular, uniform time intervals to reflects a desire for consistency and comparability in time-based analyses. Here's a breakdown of why uniform time intervals are advantageous:

Our findings, as illustrated in Fig. 4, show clear differences in how positive and negative norms develop. Positive norms tend to follow a pattern of gradual increase followed

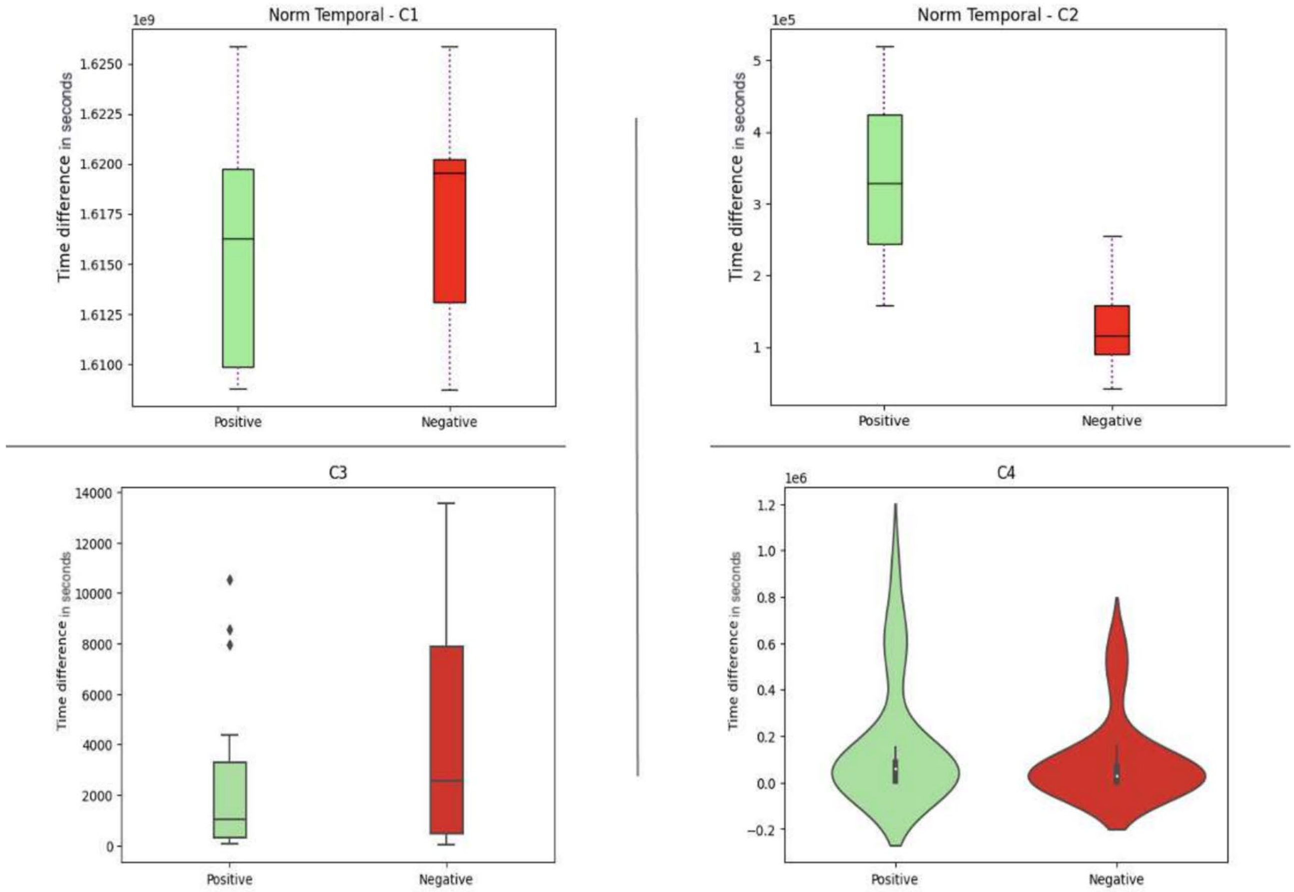

Fig. 3 Temporal analysis of positive and negative norms, where C1 is the average time from norm appearance to first supporting tweets, C2 is the average time from norm support to subsequent tweet distribution, C3 is the time gap between consecutive distributors, and C4 is the time from first support to last distributor action

by a slow decrease throughout their life cycle. Negative norms on the other hand experience a quick and significant drop in engagement within a short time frame.

The prolonged rise and subsequent fall of engagement in positive norm-associated networks suggest that users slowly adopt positive norms, and once adoption reaches its peak, engagement eventually wanes. Conversely, the rapid decline of negative norms indicates that they quickly lose relevance or influence. These contrasting patterns highlight the unique ways in which positive and negative norms emerge, spread, and ultimately fade away, assisting strategies for either preventing or promoting norms of interest.

4.4 Macro and micro analysis

We now shift the discussion to examining the intricate dynamics of norm propagation networks at the macro and micro level.

At the macro level, we observe the sequential progression of norms from their inception in original tweets (layer 1) to their broader dissemination through retweets (layer 2). This perspective allows us to understand how norms spread to wider audiences and identify key influencers in the diffusion process. Meanwhile, the micro-level networks capture more intimate user interactions through replies (layer 3) and mentions (layer 4). This multi-layered approach provides a comprehensive view of norm diffusion, encompassing both broad dissemination patterns and nuanced user interactions. Figure 5 illustrates the temporal evolution of these layers in typical networks.

Our findings reveal distinct patterns in layer distribution within these networks. For positive norms, approximately 40% of nodes occupy the second layer by the end of propagation. In contrast, negative norms see about 50% of nodes in this layer, suggesting faster initial spread compared to their positive counterparts. This is a valuable insight for detecting negative norms in online networks.

A striking difference emerges in the final stages of propagation: the proportion of layers in positive norm networks is roughly double that of negative norms. This indicates more persistent engagement with positive norms even as

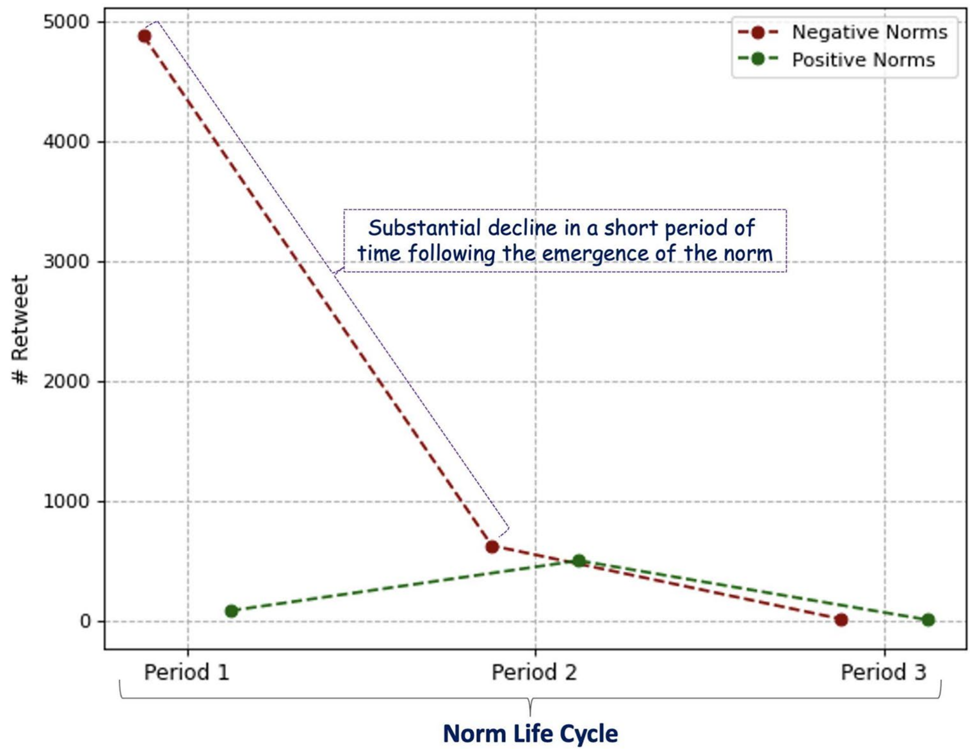

Fig. 4 The propagation speed of positive and negative norms throughout their life cycle. Positive norms experience a steady increase until reaching peak engagement, then decline over time. Negative norms display a dramatic decline in a short period of time following emergence

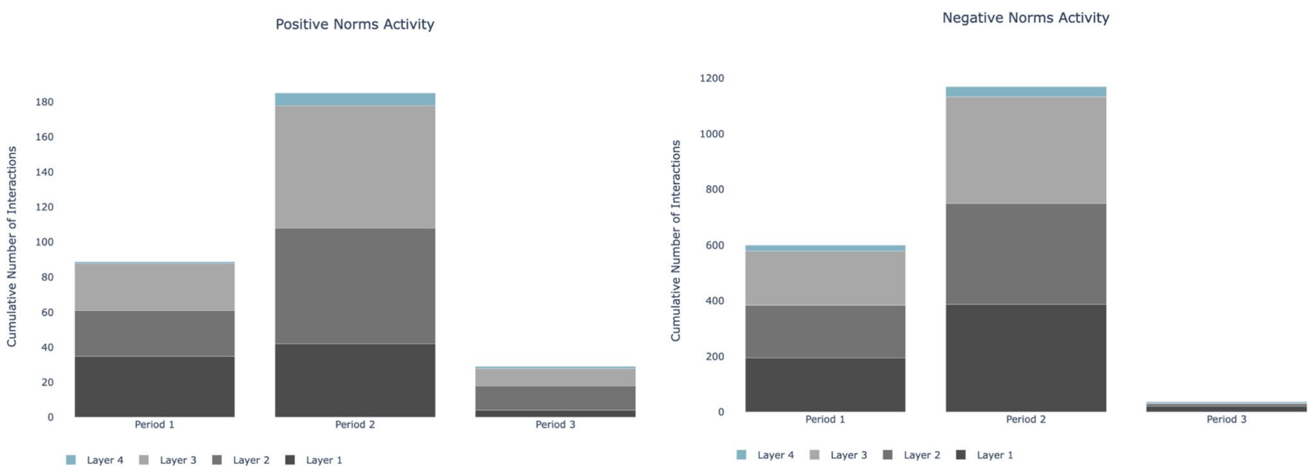

Fig. 5 Typical networks exhibit different layer sizes as a function of time. The *x* axis is the time period in a norm's life cycle. The *y* axis is the cumulative number of postings at different layers. Different colors represent different layers

they approach obsolescence, while negative norms experience a steeper decline in interaction.

To further understand user engagement, we analyzed the relationship between"likes"and sharing behaviors (retweets) over time. Positive norms maintain consistent"like"activity, reflecting sustained user interest. In contrast, negative norms see a rapid decline in "likes,"indicating briefer periods of engagement and diminishing support. These results are displayed in Fig. 6.

4.5 Community analysis of propagation networks

Community analysis serves as a crucial tool in identifying clusters of users actively engaged in norm-related discussions, sharing, and support. Particularly valuable in the context of norms, it allows us to pinpoint influential users and communities that play key roles in the adoption and propagation of these norms. By examining the social context and dynamics surrounding both negative and positive

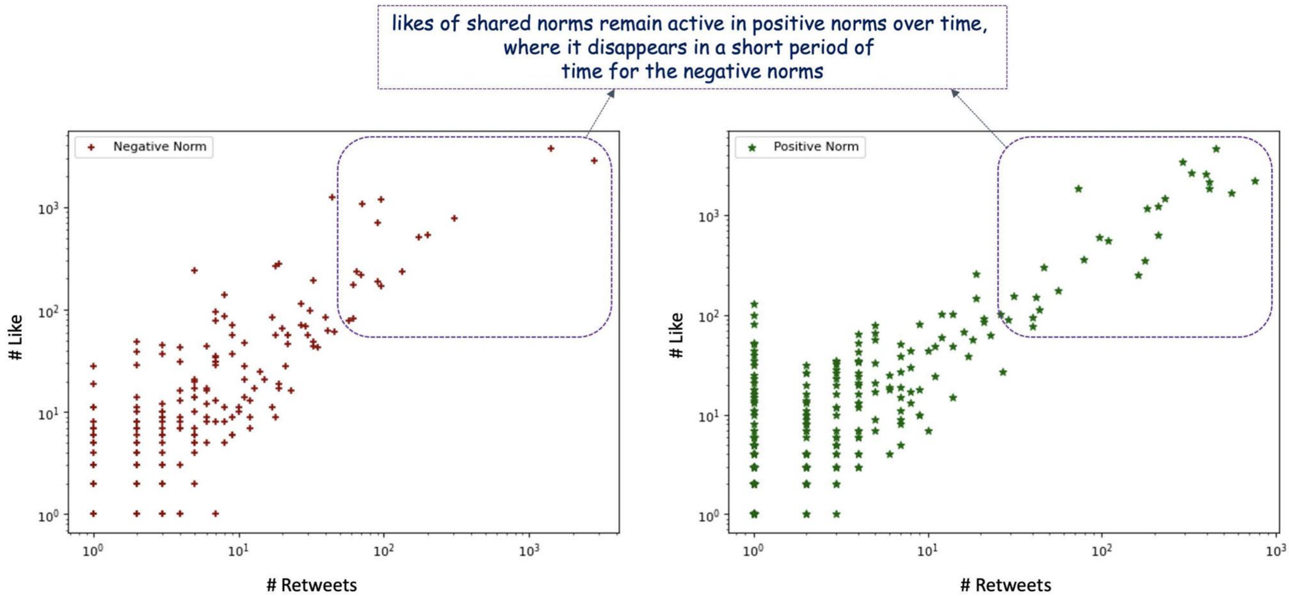

Fig. 6 The impact of“likes”on sharing behaviors of descriptive norms

norms, we gain deeper insights into the mechanisms of norm diffusion and social influence.

These findings further distinguish the distinct trajectories and engagement patterns of positive and negative norms in online social networks. They underscore the enduring influence of positive normative behaviors and offer valuable insights for fostering constructive online environments.

4.5.1 Community detection

The constructed communication network consists of edges connecting users through retweet interactions. These interactions serve as the building blocks for connections within the network. To quantify the strength of these connections, we assign weights to the edges based on how frequently one user retweets another user. This weighting scheme provides a measure of the intensity or frequency of user interactions, allowing us to capture the dynamics of communication patterns within the network.

Among the community detection algorithms we employ (Girvan and Newman 2002; Rosvall and Bergstrom 2008), Louvain is a popular choice due to its ability to optimize network modularity and efficiently identify communities (Blondel et al. 2008). The Louvain algorithm particularly stands out when it comes to handling large-scale networks, balancing computational speed and accuracy. By leveraging its optimization approach, we can extract cohesive communities and gain insights into the underlying social dynamics that shape the spread of norms.

4.5.2 Community-influencers analysis

Figure 7 illustrates the mean degree centrality analysis within communities, revealing significant differences between negative and positive norm-associated groups. Communities linked to positive norms exhibit higher centralization, suggesting the presence of key influential users who drive community growth and expansion. This finding aligns with our earlier observations, where users supporting positive norms often rely on credible sources, including healthcare professionals and governmental agencies. In contrast, communities associated with negative norms show lower centralization, indicating fewer influential users within these groups.

These distinctions are crucial for understanding the factors that contribute to norm diffusion and influence within social networks. A possible explanation for this difference is that well-connected users supporting positive norms play a pivotal role in norm propagation. These influencers engage with a broader range of users who share similar norms within their communities, allowing them to exert influence and establish connections across a wider audience.

By understanding these dynamics, we gain valuable insights into how different types of norms spread and maintain influence within online social groups. This knowledge can inform strategies for promoting positive norms and mitigating the spread of negative ones in digital spaces.

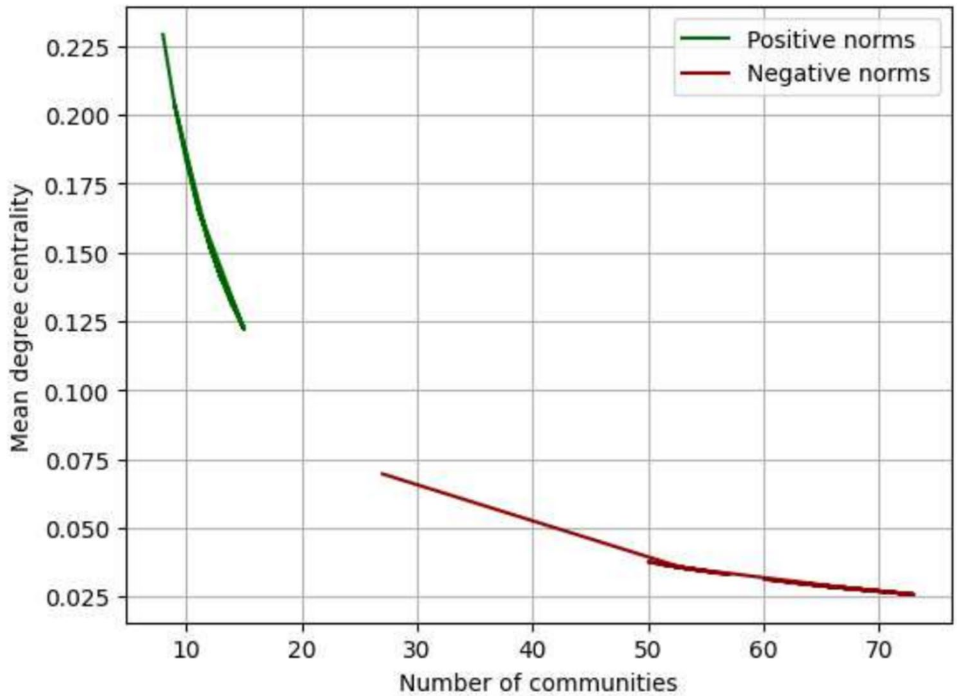

Fig. 7 The mean degree of centrality of different communities, associated with either positive or negative norms. Observe that the former are more centralized, indicating that such communities extend their reach via influencer nodes. In the context of COVID-19, this could largely be due to individuals receiving information concerning positive norms from healthcare professionals

4.5.3 Overlapping community analysis

Understanding the overlapping influence of users across various communities in a social network is essential for grasping the intricate dynamics of online interactions. We now shift the focus to how different communities interconnect and influence each other, shedding light on the information diffusion patterns of norms.

To investigate similarities in the influence of users across different communities associated with positive and negative norms, we develop a simple yet effective algorithm:

(1) Data loading: We separate the data into two distinct sets, one for positive norms and another for negative norms.
(2) Graph construction: We construct a directed graph to represent user relationships, one for each set.
(3) Node embeddings: We use Node2Vec to generate powerful continuous latent space representations for nodes in the graph, capturing users' structural roles within their communities.
(4) Influencer comparison: We identify top influencers in each community based on follower count and compare their embeddings using cosine similarity.

We then create a heatmap to visualize the overlapping influence among the largest communities, highlighting similarities in user influence. In Fig. 8, the heatmap displays similarity scores between community pairs, with higher scores indicating more similar user influence between communities.

Our findings reveal a notable distinction between positive and negative norm communities. The largest communities associated with positive norms show fewer overlapping influencers while communities linked to negative norms demonstrate more similar influencer structure across groups. This suggests that groups promoting negative norms are often guided by a cohort biased toward a particular ideology. It also suggests that influencers promoting negative norms tend to reach similar audiences whereas communities on the receiving end of positive norms tend to be more structurally diverse.

5 Conclusion

Descriptive positive and negative norms play a significant role in shaping individuals' perspectives and behaviors. Therefore, understanding them and their dynamics is crucial for designing successful promotion campaigns intended to weaken negative norms and encourage positive norms. However, conducting in-depth analysis related to community characteristics and the diffusion of descriptive norms is a complicated process influenced by many time-dependent variables. To fill existing gaps of knowledge, this paper presents a comprehensive analysis of norm propagation in online communities. The study includes macro and micro analyses, structural and temporal analyses, and activity-based norm life cycle analyses. Community analysis of propagation networks and examination of community overlapping influencers is also conducted. We shed new light on the dynamics of norms, including their underlying influence

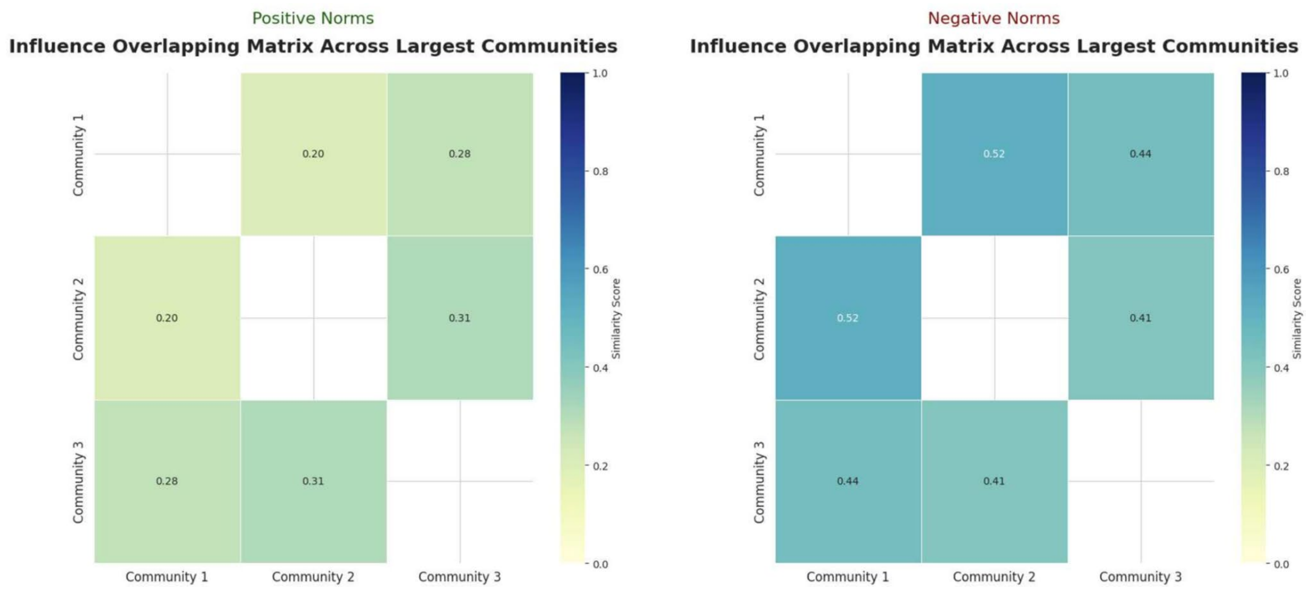

Fig. 8 A heat map of overlapping influence across the largest communities

patterns within communities, and the role of influential users and clusters in their adoption and propagation. Combining insights across analyses, this paper could further indicate that negative norms must be spread quickly in order to survive, and counteracting misinformation and disinformation could be as easy as flooding the network with the truth. This paper sheds light on the social context and dynamics of both positive and negative descriptive norms, deepening our understanding of norm diffusion and social influence processes.

Acknowledgements This work was supported in part by the National Science Foundation Program under awards No. 2323794 and No. 2323795. This paper is an extension of an earlier version presented at the 2023 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM).

Author contributions Dr. Yan Wang, Dr. Wenwen Dou prepare the dataset, Dr. Raed and Youval perform the experiment Dr. Thai clears our contribution, revise the paper. All of the authors proofread the paper.

Data availibility No datasets were generated or analysed during the current study.

Declarations

Conflict of interest The authors have no Conflict of interest to declare that are relevant to the content of this article.

References

Allen J IV, Mahumane A, Riddell J IV, Rosenblat T, Yang D, Yu H (2021) Correcting perceived social distancing norms to combat covid-19. Technical report, National Bureau of Economic Research

Blondel VD, Guillaume JL, Lambiotte R, Lefebvre E (2008) Fast unfolding of communities in large networks. J Statist Mech: Theory Exp 2008(10):P10008

Cao J, Liu F, Shang M, Zhou X (2021) Toward street vending in post covid-19 China: social networking services information overload and switching intention. Technol Soc 66:101669

Chen Y, Hong Y (2015) Different ways to resolve discrepancy between descriptive and injunctive norms across cultures. J Cross Cult Psychol 46(10):1316–1319

Dillard JP, Tian X, Cruz SM, Smith RA, Shen L (2021) Persuasive messages, social norms, and reactance: a study of masking behavior during a covid-19 campus health campaign. Health communication, pp 1–11

Dow BJ, Johnson AL, Wang CS, Whitson J, Menon T (2021) The COVID-19 pandemic and the search for structure: social media and conspiracy theories. Soc Personal Psychol Compass 15(9):e12636

Eskytė I, Lawson A, Orchard M, Andrews E (2020) Out on the streets-crisis, opportunity and disabled people in the era of covid-19: reflections from the UK. Alter 14(4):329–336

Gadzekpo A, Tietaah GKM, Yeboah-Banin AA, Kwame Ampofo Adjei D (2023) Media coverage of covid-19 vaccines: sources of information, and verification practices of journalists in Ghana. J Commun Healthcare, pp 1–15

Gao S, Wang Y, Platt L (2021) Modeling us health agencies' message dissemination on twitter and users' exposure to vaccine-related misinformation using system dynamics. In: ISCRAM 2021 conference proceedings–18th international conference on information systems for crisis response and management

Gao S, Wang Y, Thai MT (2023) Investigating the dynamics of social norm emergence within online communities. arXiv preprint arXiv:2301.00453

Girvan M, Newman MEJ (2002) Community structure in social and biological networks. Proc Natl Acad Sci 99(12):7821–7826

Hassell HJG, Wyler EE (2019) Negative descriptive social norms and political action: people aren't acting, so you should. Political Behav 41:231–256

Hawkins RXD, Goodman ND, Goldstone RL (2019) The emergence of social norms and conventions. Trends Cogn Sci 23(2):158–169

Himelboim I, Borah P, Lee DKL, Lee J, Su Y, Vishnevskaya A, Xiao X (2023) What do 5g networks, bill gates, agenda 21, and Ganon have in common? sources, distribution, and characteristics. New Media Soc, pp 14614448221142800

Hu S, Leung H-F (2017) Local norm phenomena in multi-agent systems under community networks. In: Proceedings of the 16th conference on autonomous agents and MultiAgent Systems, pp 1570–1572

Latkin CA, Dayton L, Kaufman MR, Schneider KE, Strickland JC, Konstantopoulos A (2022) Social norms and prevention behaviors in the United States early in the covid-19 pandemic. Psychol Health Med 27(1):162–177

McDonald RI, Fielding KS, Louis WR (2014) Conflicting social norms and community conservation compliance. J Nat Conserv 22(3):212–216

Neville FG, Templeton A, Smith JR, Louis WR (2021) Social norms, social identities and the covid-19 pandemic: theory and recommendations. Soc Personal Psychol Compass 15(5):e12596

Romer D, Jamieson KH (2020) Conspiracy theories as barriers to controlling the spread of covid-19 in the US. Soc Sci Med 263:113356

Rosvall M, Bergstrom CT (2008) Maps of random walks on complex networks reveal community structure. Proc Natl Acad Sci 105(4):1118–1123

Saint SA, Moscovitch DA (2021) Effects of mask-wearing on social anxiety: an exploratory review. Anxiety Stress Coping 34(5):487–502

Savarimuthu BTR, Cranefield S (2011) Norm creation, spreading and emergence-a survey of multi-agent based simulation models of norms. J Multiagent Grid Syst 7(1):21–54

Shu K, Mahudeswaran D, Wang S, Liu H (2020) Hierarchical propagation networks for fake news detection: investigation and exploitation. In: Proceedings of the nternational AAAI conference on web and social media

Sikali K (2020) The dangers of social distancing: How covid-19 can reshape our social experience. J Commun Psychol

Varma JK, Taylor J, Sharfstein JM (2023) Planning for the next pandemic: lab systems need policy shift to speed emerging infectious disease warning and tracking: article examines responses of us laboratories to covid-19 and suggests improvements to speed disease warning and tracking. Health Affairs 42(3):366–373

Vosoughi S, Roy D, Aral S (2018) The spread of true and false news online. Science 359(6380):1146–1151

Wang D, Jiang M, Syed M, Conway O, Juneja V, Subramanian S, Chawla NV (2020) Calendar graph neural networks for modeling time structures in spatiotemporal user behaviors. In: Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining, pp 2581–2589

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