

Network Structure Impacts the Synchronization of Collective Beliefs

Madalina Vlasceanu

Postdoctoral Fellow, Department of Psychology at New York University,
New York, NY, USA

Corresponding author
mov@princeton.edu

Michael J. Morais

Graduate Student, Princeton Neuroscience Institute, Princeton University,
Princeton, NJ, USA

Center for Statistics and Machine Learning, Princeton University,
Princeton, NJ, USA
mjmorais@princeton.edu

Alin Coman

Associate Professor, Department of Psychology, Princeton University,
Princeton, NJ, USA
Princeton School of Public and International Affairs, Princeton University,
Princeton, NJ, USA
acoman@princeton.edu

Abstract

People's beliefs are influenced by interactions within their communities. The propagation of this influence through conversational social networks should impact the degree to which community members synchronize their beliefs. To investigate, we recruited a sample of 140 participants and constructed fourteen 10-member communities. Participants first rated the accuracy of a set of statements (pre-test) and were then provided with relevant evidence about them. Then, participants discussed the statements in a series of conversational interactions, following pre-determined network structures (clustered/non-clustered). Finally, they rated the accuracy of the statements again (post-test). The results show that belief synchronization, measuring the increase

in belief similarity among individuals within a community from pre-test to post-test, is influenced by the community's conversational network structure. This synchronization is circumscribed by a degree of separation effect and is equivalent in the clustered and non-clustered networks. We also find that conversational content predicts belief change from pre-test to post-test.

Keywords

collective beliefs – belief change – social networks – misinformation

1 Introduction

Human societies are characterized by extensive communicative exchanges. This dynamic information flow has been shown to exert a strong influence on people, impacting their individual memories (Cuc, Koppel, & Hirst, 2007), their beliefs (Vlasceanu & Coman, 2020), and their behaviors (Frankel & Swanson, 2002). It has also been found to affect collective-level phenomena, such as the formation of collective memory (Coman, Momennejad, Drach, Geana, 2016), collective beliefs (Vlasceanu, Morais, Duker, Coman, 2020), and collective decision-making (Bahrami et al., 2012). A growing body of work has been focusing on the cognitive and social processes involved in these collective phenomena (Vlasceanu, Enz, Coman, 2018; Borge, Ong, Rosé, 2018), revealing the importance of network structure in their emergence. Prior work investigating network influences on collective beliefs has mainly relied on theoretical (Friedkin & Johnsen, 1990) and simulation-based approaches (Friedkin, Proskurnikov, Tempo, & Parsegov, 2016), revealing, for instance, how beliefs can converge in social networks (Friedkin et al., 2016). Here, we take an empirical approach to revealing processes underlying collective belief change. While a belief is defined as a statement held to be true (Schwitzgebel, 2010), a collective belief is characterized by a group of individuals' joint commitment towards a particular belief (Gilbert, 2000; Bouvier, 2004). A central feature of beliefs is their dynamic nature, making them susceptible to change (Bendixen, 2002). Indeed, prior work on the synchronization of beliefs revealed that conversations within 3-member groups can change beliefs, leading to their coordination (Wilkes-Gibbs & Clark, 1992). Additionally, community members, after being exposed to a public speaker's beliefs and having conversations about them within their social networks, align with the public speaker's beliefs and therefore become more synchronized with each other (Vlasceanu et al., 2020). This effect, the authors showed, was driven by a memory mechanism by

which beliefs become stronger when their mnemonic accessibility is increased and weaker when their accessibility is decreased (Vlasceanu & Coman, 2018).

Here, we are interested in expanding this work and programmatically exploring how the synchronization of collective beliefs is influenced by a community's network structure. We are studying this influence considering both a time-independent topological mapping typically used in network analysis (Watts & Strogatz, 1998), as well as the temporal sequencing of conversational interactions in networks (Tang, Musolesi, Mascolo, Latora, 2009). Both modalities of experimentally manipulating networks have been found to impact collective-level phenomena. Coman and colleagues (2016) manipulated the clustering coefficient of conversational networks and found that non-clustered networks reached higher mnemonic convergence than clustered networks. To showcase the impact of temporal sequencing of conversations in driving the formation of collective memories, Momennejad, Duker, & Coman (2019) manipulated when conversations occurred between people who bridged network clusters (either early or late during the community's conversations). They found that early conversations between bridge individuals lead to increased community-wide convergence compared to late conversations between bridge individuals. No such investigation has been conducted to explore the dynamics involved in the formation of collective beliefs.

To investigate how network structure affects collective belief synchronization in social networks, we designed an experiment in which participants were invited in the lab in groups of 10, each group forming a lab-created community (Fig. 1). One hundred and forty participants enrolled in the study through Princeton University's recruitment system. First, participants rated the accuracy of a set of statements, pretested and selected for their moderate believability (belief pre-test). In reality, half of the statements were scientifically accurate, and the other half were scientifically inaccurate. They were then provided with relevant evidence supporting or refuting half of the initial statements (target items); the other half of the initial statements for which no evidence was provided were considered baseline items. The pieces of evidence were constructed to vary in a randomized fashion on two features (anecdotal/scientific, normative/non-normative) to increase the credibility of the stimuli and the external validity of the study. After reading the evidence, participants were asked to discuss the statements with each other, in a series of dyadic conversational interactions within their community. These interactions were computer-mediated and followed one of two pre-determined network structures, clustered and non-clustered (Fig. 1). Participants were assigned to their position in clusters randomly and did not have knowledge of the structure of the network. Finally, participants rated the accuracy of the initial set of statements again (belief post-test).




FIGURE 1 Network structures: clustered (Panel A), non-clustered (Panel B). Circles represent participants, and links represent conversations. Numbers in red indicate the sequence of conversations.

Our first hypothesis was that the conversational network structure would impact the synchronization of collective beliefs, measured as the increase in belief similarity of all pairs of individuals within each network from pre-test to post-test. We also hypothesized that this effect would be circumscribed by a degree of separation effect, such that individuals closer to each other in the network would become more similar than individuals further away from each other in the network. Lastly, we hypothesized that the conversational content would influence the direction and degree of belief change, such that statements endorsed in conversations would increase in believability whereas statements opposed in conversations would decrease in believability.

2 Methods

2.1 Open Science Practices

The data and materials can be found on our open science framework page: <https://osf.io/72ayt/>.

The data analysis (in python) can be viewed as a jupyter notebook: <https://github.com/mvlasceanu/CollectiveBeliefs>.

All measures, manipulations, and exclusions are disclosed.

2.1.1 Participants

A total of 140 Princeton undergraduate students (63% women; $M_{age} = 19.47$, $SD_{age} = 1.53$) were recruited for the study. They participated in the study for either monetary compensation or research credit. Participants were grouped into fourteen 10-member networks. The statistical power afforded by this sample size was deemed adequate given effect sizes obtained in previous studies using a similar sample size and experimental paradigm (Coman et al., 2016; Vlasceanu et al., 2020). We note that the aggregation procedure used to compute network-wide belief synchronization reduces the standard errors around

the mean and results in a highly accurate estimate of the true value. The study was approved by the Institutional Review Board at Princeton University, and was conducted in accordance with IRB guidelines and regulations. Informed consent was obtained from all participants.

2.1.2 Stimulus Materials

A set of 16 statements of moderate believability were selected from a larger set of 32 statements used in a study by Vlasceanu and Coman (2020) (e.g., *"Eating carrots will make eyesight sharper"*). This study, conducted on a sample of Princeton undergraduate students ($N = 200$; $M_{age} = 19.49$, $SD_{age} = 1.39$; 64% women), collected believability ratings for each statement (i.e., *"How accurate or inaccurate do you think this statement is"*) on a scale from 0-*Extremely Inaccurate* to 100-*Extremely Accurate*). The final set of 16 statements we used in the current study were selected such that for each of them, their level of believability in this population of interest was in the moderate range ($M = 51.41$, $SD = 4.18$, on a 0 to 100-point scale). Even though each of these 16 statements was equally believable by design, half of them were scientifically accurate, while the other half were scientifically inaccurate, as determined by published scientific papers.

We constructed a set of 8 pieces of direct evidence, either in favor or against half of the initial 16 statements. The statements for which no evidence was presented were considered baseline items. The pieces of evidence were constructed such that they argued in favor of the initial statement if the statement was accurate (e.g., *"Children who spend less time outdoors are at greater risk to develop nearsightedness, study shows"*) and against the initial statement if the statement was inaccurate (e.g., *"Eating carrots does not make eyesight sharper, study shows"*). To increase external validity, these pieces of evidence were constructed and displayed to participants as if they were tweets collected from the Twitter platform. To increase the variability of these pieces of evidence, mirroring content frequently found on Twitter (Zhang et al., 2019), we counterbalanced the phrasing suggesting the evidence posted is the result of a study, with phrasing suggesting the evidence posted is anecdotal, such that each statement was in either counterbalancing condition with equal probability randomly assigned across participants. Also consistent with content on Twitter (Kim, 2018), we counterbalanced the number of retweets each of these tweets had, such that each statement either had a large or a small number of retweets also with equal probability randomly assigned across participants. Moreover, for each piece of evidence, the sources were constructed to be as similar as possible while allowing some variability, to maintain the credibility of the stimuli. In each case, the person tweeting was depicted as a white middle-aged male, with a common name and appearance. The dates of the tweets were randomly

chosen from dates in the month of July 2019. Even though these varying features of the evidence varied randomly across participants and items, they were held constant within a network.

2.1.3 Design and Procedure

The 140 participants were split in 14 lab-created communities of 10 participants. Each community was assembled separately and was comprised of individuals who arrived in the lab at the same time. These communities were randomly assigned to either the Clustered network structure condition (7 networks) or the Non-Clustered network structure condition (7 networks), following procedures by Coman and colleagues (2016). These network structures resemble real-world small networks characterized by varying degrees of clustering (Watts & Strogatz, 1998; Tang et al., 2009). Once assigned to condition, participants went through four experimental phases. Participants were told they would participate in an experiment about people's evaluation of information and were directed to the survey on the Qualtrics platform. After completing the informed consent form, participants were directed to the first phase (pre-test), in which they rated a set of 16 statements (one on each page) by indicating the degree to which they believed each statement (i.e., "*How accurate do you think this statement is,*" from 1-*Extremely inaccurate* to 100-*Extremely accurate*). Then, in the second phase (evidence phase), participants were exposed to a sub-set of 8 pieces of evidence (in the form of tweets, one on each page), half of which argued in favor and the other half argued against the initial statements (i.e., target statements). Eight statements presented initially constituted baseline items. Each of the 16 statements were randomly assigned to either a target or a baseline status. To ensure participants processed the evidence information, they were instructed to rate each tweet on how convincing, rigorous, widespread, and personal it appears to them, as well as how likely they would be to share it. In the third phase (conversational phase), participants were directed to another software platform (i.e., Software Platform for Human Interaction Experiments; SoPHIE) that allows fluent computer-mediated interactions among participants. At this stage, participants were instructed to discuss the information from the study with each other, in a series of dyadic conversations (each with a different partner). Conversations took the form of interactive exchanges in a chat-like computer-mediated environment in which participants typed their responses. Each participant had three conversations, each lasting 150 seconds. Participants in the Clustered condition ($n = 70$ participants; seven 10-member networks) communicated according to a network structure characterized by two subclusters (Fig. 1A), whereas participants in the Non-clustered condition ($n = 70$ participants; seven 10-member networks), communicated according to a network structure characterized by a single large

cluster (Fig. 1B). The number of participants per network ($n = 10$), the sequencing of conversational interactions, and the number of conversations each participant had (i.e., three) were kept constant between the two conditions. Finally, in the fourth phase (post-test) participants rated again the believability of the initial 16 statements, after which they were asked to complete a series of demographic information and were debriefed.

2.1.4 Dependent Variables

We operationalized *rational belief update* as the belief change from pre-test to post-test in the direction corresponding to incorporating the available evidence (Table 1, eq. 1). For statements with supporting evidence – the rational update is to increase in believability from pre-test to post-test. For statements followed by refuting evidence – the rational update is to decrease in believability from pre-test to post-test. Through counterbalancing, we ensured that participants could not trivially infer that “correct” updates must be in one direction. We also operationalized *belief similarity* as the correlation of the beliefs held by a pair of individuals at a given timepoint (Table 1, eq. 2), and *belief synchronization* as the increase from pre-test to post-test in the average of all belief similarity scores within a network (i.e., all correlations of beliefs of all the possible pairs of individuals; table 1, eq. 3). The *degree of separation* denoted the minimum number of links a given pair of nodes within a network

MEASURE/DEFINITION	FORMULA	FIGURE
RATIONAL BELIEF UPDATE The belief change from pre-test to post-test in the direction corresponding to incorporating the available evidence.	$\Delta B_r = E \square \Delta B \quad (1)$ Evidence = {Refuting, Supporting} / $E = \{-1, 1\}$ ΔB = change in belief from pre to post	
BELIEF SIMILARITY The correlation of the beliefs held by a pair of individuals at a given timepoint	$BSim_{ij} = r_s(B_{p_i}, B_{p_j}) \quad (2)$ B_{p_i} = belief vector of participant i	
BELIEF SYNCHRONIZATION The increase from pre-test to post-test in the average of all belief similarity scores within a network	$BSynch = \overline{r_s^{POST}} - \overline{r_s^{PRE}} \quad (3)$ $\overline{r_s} = \frac{1}{\binom{N}{2}} \sum_{i < j} r_s(B_{p_i}, B_{p_j}) \quad (4)$ $\overline{r_s}$ = average belief similarity in a network N = number of nodes (participants in a network) $\binom{N}{2}$ = number of pairs of nodes	
CLUSTERING COEFFICIENT (STATIC) Measure of the degree to which nodes in a graph tend to cluster together	$\overline{C} = \frac{1}{N} \sum_i \frac{\# \text{ of edges in } G^{N_i}}{\binom{N_i}{2}} \quad (5)$ N_i = neighbors of node i , G = network, G^{N_i} = neighbor subnetwork of node i k_i = # of neighbors of node i	
CLUSTERING COEFFICIENT (TEMPORAL) Measure of the degree to which nodes in a graph tend to cluster together at each moment in time.	$\overline{C_{0,T}} = \frac{1}{N} \sum_i \frac{\sum_{t=0}^T \# \text{ of edges in } G_i^{N_i}(0, T)}{T \binom{k_i(T)}{2}} \quad (6)$ T = conversational round	

TABLE 1 Definitions, equations, and figures for the dependent variables

are away from each other (i.e., how many conversations would need to occur for information from one node to reach the other).

2.1.5 Conversational Coding

We coded the conversations' content for *recall* and for *belief endorsement*. The recall coding involved a binary system in which a statement was labeled as either mentioned or not mentioned by each participant in each of their conversations. Thus, a participant could have mentioned a given statement anywhere from 0 times up to 3 times (once in each of their 3 conversations). For each participant, their conversational partners could have also mentioned a particular statement up to 3 times. Thus, for each participant, the *joint recall* measure we used denotes the number of times either the participant or their interaction partners mentioned each statement (i.e., from 0 to 6). Coding for belief endorsement in conversations entailed the additional factor of valence, which denotes whether a mentioned statement is being endorsed or refuted in conversation. Thus, the belief coding involved a larger interval (i.e., from -3 to +3), in which a mentioned statement was labeled as either strongly endorsed (+3), moderately endorsed (+2), slightly endorsed (+1), just mentioned (0), slightly opposed (-1), moderately opposed (-2), or strongly opposed (-3). Again, for each participant we accounted for both their own input and their 3 conversational partners' inputs for each statement, to form a measure of *joint belief* (i.e., from -18 to +18). To make this interval comparable and consistent to the recall interval, we normalized it by dividing it by 3. In each case, ten percent of the data were double coded for reliability (Cohen $\kappa > 0.89$).

3 Results

Since there was no difference in the initial level of believability of scientifically accurate statements ($M = 51.36$, $SD = 12.68$) and scientifically inaccurate statements ($M = 51.93$, $SD = 12.61$), $p = 0.66599$, which were pre-tested to be as similar as possible, we combined them for the rest of the analyses conducted. Similarly, we did not find any differences in rational updating between the varying features of the pieces of evidence (e.g., anecdotal/scientific) in a repeated-measures ANOVA with rational belief update as the dependent variable, and evidence type as a within-subject variable $F(3, 417) = 1.09$, $p = 0.353$. Therefore, we combined them for the rest of the analyses.

To investigate whether individuals' beliefs became more similar post conversations relative to pre-conversations, we ran a linear mixed model (using R version 3.1.0, with lme4 version 1.1.21) with belief similarity (for the target

items) as the dependent variable, time-point (pre-test vs. post-test) as the fixed effect, and by-network random intercepts which nests the data by networks, and found that belief similarity scores (i.e., pairwise belief correlations) at post-test ($M = 0.0649$, $SD = 0.380$) were significantly higher ($\beta = 0.06$, $SE = 0.01$, $t(139) = 3.5$, $p < 0.001$) than at pre-test ($M = 0.0079$, $SD = 0.374$) (Fig. 3A).

Collective belief synchronization is dependent on network structure. To investigate our first hypothesis, that network structure would impact belief synchronization, we conducted two sets of analyses. First, we compared the levels of synchronization achieved by the two network structures employed. We found that synchronization was equivalent in the Clustered ($M = 0.0257$, $SD = 0.090$) compared to the Non-Clustered ($M = 0.0254$, $SD = 0.067$) condition (Fig. 3B), in a linear mixed model with belief similarity difference (for the target items) as the dependent variable, condition (clustered vs. non-clustered) as the fixed effect, and by-network random intercepts ($\beta = 0.02$, $SE = 0.01$, $t(13) = 1.52$, $p = 0.15$).

A more complex analysis that takes into account the degree of separation among community members involved the construction of a hypothesis matrix for each of the 2 network structures used (clustered and non-clustered; Fig. 2A), following methods by Coman and colleagues (2016). We reasoned that participants who had a conversation with one another would synchronize their beliefs more than participants who were two degrees away from one another in the network, and so on. In this context, a hypothesis matrix represents the hypothesized distance in belief ratings between any two participants in a community, depending on their distance in the network. In constructing these hypotheses matrices, we used the range of the entire empirical belief synchronization scores (i.e., post minus pre scores of the target items belief correlations of each pair of participants within a network) collected in each condition. We split the intervals of these distributions into quintiles (clustered) or terciles (non-clustered), and mapped the boundary values of these intervals to the corresponding degree of separation in the participant by participant matrix (e.g., first quantile boundary value was mapped on the first degree of separation, second quantile boundary value was mapped on the second degree of separation, etc.). This density matching procedure was employed such that the hypotheses matrices can be comparable to the corresponding empirical matrices. The empirical data matrices were constructed by allocating the observed synchronization scores of each pair of participants to the corresponding cells in the participant by participant matrices.

Once we constructed both the hypothesis matrices and the empirical matrices (Fig. 2A), we ran a non-parametric statistical test to assess our first hypothesis – that network structure would circumscribe belief synchronization. Specifically, we tested whether the empirical matrices are better explained




FIGURE 2 Panel A. Hypothesis Matrix of Clustered (top left) and Non-Clustered (bottom left) Network Structures. Empirical Data of Clustered (top right) and Non-Clustered (bottom right) Network Structures. Panel B. Histogram of belief similarity correspondence between empirical data (Clustered on the top, Non-Clustered on the bottom) and 10,000 bootstrapped random hypothesis matrices matched to the experiments' network features (10 nodes and 3 ties per node). The vertical lines represent the correspondence of the hypothesis matrices with the true network structures.

by their corresponding hypothesis matrix than by any other hypothesis matrices that could exist based on the same constraints of building the conversational network. To do so, we simulated 10,000 bootstrapped random hypothesis matrices matched to the true hypothesis matrices' features (10 nodes, 3 ties per node) using the same data mapping procedure used for the true hypothesis matrices. We then compared the empirical matrices to both the true hypothesis matrices and to each of these 10,000 random hypothesis matrices by running a series of quantile correlations. We found that the empirical data was better explained by the true hypothesis matrices (the true network structures) than by most of the simulated hypothesis matrices for both the Clustered ($p = 0.0462$) and the Non-Clustered ($p = 0.0075$) conditions (Fig. 2B). The p-values in this non-parametric test represent the probability that any of these simulated hypothesis matrices better explain the data than the true ones. Thus, we found support for our first hypothesis, that belief synchronization is determined by the structure of the network individuals are embedded in.

Further support for the hypothesis that the level of belief synchronization exhibits a degree of separation effect, comes from a repeated measures ANOVA nested by network, with belief similarity difference as the dependent variable and degree of separation (1–5) as the independent within-subject variable, showing a main effect of degree of separation $F(4, 52) = 4.21, p < 0.004, \eta_p^2 = 0.24$. Moreover, in a linear mixed model with belief similarity difference (for the target items) as the dependent variable, degree of separation (with 5 levels) as the fixed effect, and by-network random intercepts, we find that, as hypothesized, individuals situated 1-degree of separation away ($\beta = 0.057, SE = 0.017, t(48) = 3.39, p < 0.0014$) became significantly more similar to each other, whereas individuals 2-degrees away ($\beta = 0.026, SE = 0.017, t(47) = 1.54, p = 0.129$) or further did not (3-degrees away: $\beta = 0.009, SE = 0.02, t(101) = 0.46, p = 0.641$, 4-degrees away: $\beta = -0.034, SE = 0.03, t(230) = -1.13, p = 0.258$, and 5-degrees away: $\beta = 0.032, SE = 0.04, t(447) = 0.79, p = 0.427$) (Fig. 3C). This suggests that the belief influence within a community only travels one degree of separation away from the originating source.

Conversational content predicts belief change. So far, we have showed that people’s beliefs become more similar after conversations, and this synchronization occurs according to their position in conversational networks. However, we have not yet considered the content of their conversations, and how these conversations lead to belief change. Here, we are investigating our last hypothesis, that endorsing or opposing beliefs in conversations influences rational belief update. First, focusing on simple recall, we conducted a linear mixed model with rational belief update (for the target items) as the dependent variable, joint conversational recall as the fixed effect, as well as by-participant random intercepts and by-network random intercepts (Fig. 4B). We found that




FIGURE 3 **Panel A.** Belief similarity at pre-test and post-test. **Panel B.** Belief similarity change (pre-test minus post-test) in the Clustered versus Non-Clustered Networks. **Panel C.** Belief similarity change (pre-test minus post-test) by degree of separation. Error bars represent ± 1 standard errors of the mean.




FIGURE 4 Rational belief update as a function of conversational belief endorsement (Panel A) and as a function of conversational recall (Panel B). Error bars represent 95% bootstrapped confidence intervals on the means.

joint recall significantly predicts rational belief update ($\beta = 1.62$, $SE = 0.90$, $t(941) = 3.21$, $p < 0.0014$). This was true regardless of whether the speaker ($\beta = 2.91$, $SE = 0.47$, $t(892) = 3.42$, $p < 0.001$) or the listener ($\beta = 2.54$, $SE = 0.83$, $t(850) = 3.04$, $p < 0.0024$) was the one mentioning the beliefs. Second, to assess the impact of endorsement/opposition on rational belief update, we conducted a linear mixed model with rational belief update (for the target items) as the dependent variable, joint conversational belief endorsement as the fixed effect, as well as by-participant random intercepts and by-network random intercepts (Fig. 4A). We found that joint belief endorsement also significantly predicts rational belief update ($\beta = 1.92$, $SE = 0.33$, $t(1079) = 5.71$, $p < 0.001$). Again, just like for the memory effect, the belief effect was true regardless of whether the speaker ($\beta = 2.99$, $SE = 0.62$, $t(1097) = 4.78$, $p < 0.001$) or the listener ($\beta = 2.58$, $SE = 0.54$, $t(1110) = 4.93$, $p < 0.001$) was the one supporting/opposing the belief. When running these two predictors together in the same linear mixed model with rational belief update (for the target items) as the dependent variable, and both joint conversational recall *and* joint conversational belief endorsement as fixed effects, as well as by-participant random intercepts and by-network random intercepts, we found that only belief endorsement ($\beta = 1.73$, $SE = 0.35$, $t(1099) = 4.92$, $p < 0.001$) and not recall ($\beta = 0.92$, $SE = 0.49$, $t(950) = 1.87$, $p = 0.06$) significantly predicts rational belief update. This suggests that both mentioning a statement and qualifying the level of endorsement towards it seem to explain rational belief update; when taken together most of the variance in belief update is explained by the latter variable.

4 Discussion

Human societies are organized in social networks of interconnected individuals that exchange information. Here, we show that when individuals have conversations within their communities, they synchronize their beliefs according to the conversational structure of the network they are embedded in, such that the closer one is in the network to another individual the more similar their beliefs become.

These findings complement and extend prior research in meaningful ways. First, the findings reported here are consistent with prior work on the formation of collective memory, extending it to the formation of collective beliefs. While Coman and colleagues (2016) showed that network structure influences a community's collective memory, here, we show that network structure influences a community's collective beliefs. Moreover, in both studies this influence was explained by a degree of separation effect, by which the smaller the degree of separation between any two participants, the more similar their beliefs became.

Second, these results align with past work on the formation of collective beliefs showing that following conversational interactions, people's beliefs become more coordinated (Wilkes-Gibbs & Clark, 1992), and synchronized (Vlasceanu et al., 2020). Our findings extend these effects and in addition to showing an increased overall synchronization following conversations, we show how the network structure of the conversations determines the degree of synchronization. Moreover, in addition to prior work, we show that not only do conversations increase belief similarity among communities, the content of the conversations is also important in determining the direction and degree of belief change. The more a statement is endorsed in conversation, the more its endorsement increases post conversation for both conversational partners. Similarly, the more a statement is opposed in conversation, the more its endorsement decreases post-conversation for the conversational partners.

And third, these findings are also consistent with work suggesting that even though there are 6 degrees of separation in real-world networks (Milgram, 1967), influence only spreads 3 degrees away from the originating source (Fowler & Christakis, 2010). In the domain of beliefs, it seems that the degree of influence is even shorter, with belief influence only propagating one degree away from the source. This suggests that the propagation of beliefs might involve different processes from those prompting the propagation of other psychologically grounded phenomena such as memory (Coman et al., 2016), altruism (Fowler & Christakis, 2010), and loneliness (Cacioppo, Fowler, Christakis, 2009).

Moreover, in addition to a topological mapping typically used in classic studies involving static networks (Watts & Strogatz, 1998; table 1, eq. 5), here,

we over imposed a temporal network framework, to gain additional ecological validity (Table 1, eq. 6). This distinction between the topological and temporal mapping of networks is important from a theoretical perspective – while static networks aggregate information flowing through a network over time, dynamic networks account for the temporal dimension of real-world networks (Tang et al., 2009). This difference between static and dynamic networks can also explain the equivalent level of belief synchronization produced by the two network structures used in our experiment: clustered and non-clustered. These two network structures were chosen to be as distinct as possible in terms of clustering from a topological perspective (conditioned on maintaining the constraint of regular networks with 10 nodes and 3 links per node). When only taking into account the topological mapping of the network structure, the clustered network's global clustering coefficient is 0.4, and the non-clustered network's global clustering coefficient is 0 (Freeman, 1978; table 1, eq. 5). However, the two network structures tested here were too similar in terms of clustering coefficient when taking into account the temporal dimension. When considering the temporal aspect of the networks, the clustered network's temporal clustering coefficient is 0.066, and the non-clustered network's temporal clustering coefficient is 0 (Tang et al., 2009; table 1, eq. 6). Therefore, these two temporal structures are too similar to each other to render observable differences in belief synchronization, at a small sample size. We posit that increasing the sample size or increasing the difference in temporal clustering coefficient between the network structures, would indeed render meaningful differences in belief synchronization, as suggested by our main finding that network structure plays an important role in the way people align their beliefs.

Finally, this work adds to the literature on social influence in networks as well as the literature on evolved social cultural learning (Richerson & Boyd, 2005; Mesoudi, 2005) by extending theoretical (Friedkin & Johnsen, 1990) and simulation-based (Friedkin et al., 2016) approaches to investigating collective beliefs. For instance, Friedkin and colleagues algorithmically showed how the U.S. population changed their views on whether the Iraq invasion by the United States was justified (Friedkin et al., 2016). Here, we extend such efforts with an empirical investigation.

The current investigation opens several avenues for future research. First, an important extension of this work involves programmatically investigating the impact of other network features on belief synchronization. What are the temporal and topological features of the network that results in the highest degree of synchronization? Second, another noteworthy trajectory could involve assessing the impact of trust on belief synchronization. Would manipulating trust among conversational partners (e.g., high vs. low) before individuals engage in conversational interactions differentially impact belief

synchronization? Third, investigating the effects of belief synchronization on behavior would be a research trajectory of critical importance especially in the current socio-political climate. What is the impact of collective belief change on communities' behaviors such as voting or living sustainably? And lastly, exploring the synchronization of worldview-type beliefs such as ideological or religious beliefs would meaningfully extend the current investigation. Do beliefs held with more conviction follow the same pattern of synchronization in a community?

Understanding the mechanisms by which collective beliefs take shape and change over time is essential from a theoretical perspective (Vlasceanu, Enz, Coman, 2018), but perhaps even more urgent from an applied point of view. This urgency is fueled by recent findings showing that false news diffuse farther, faster, deeper, and more broadly than true ones in social networks (Vosoughi, Roy, Aral, 2018), and that news can determine what people discuss and even change their beliefs (King, Schneer, White, 2017). And given that beliefs influence people's behaviors (Shariff & Rhemtulla, 2012; Mangels, Butterfield, Lamb, Good, Dweck, 2006; Ajzen, 1991; Hochbaum, 1958), understanding the dynamics of collective belief formation is of vital social importance as they have the potential to affect some of the most impending threats our society is facing from pandemics (Pennycook, McPhetres, Zhang, Rand, 2020) to climate change (Benegal & Scruggs, 2018). Thus, policy makers could use such findings in designing misinformation reduction campaigns targeting communities (Dovidio & Esses, 2007; Lewandowsky et al., 2012). For instance, these findings suggest such campaigns be sensitive of the conversational network structures of their targeted communities. Knowing how members of these communities are connected, and leveraging the finding that people synchronize their beliefs mainly with individuals they are directly connected to, could inform intervention designers how communities with different connectivity structures might respond to their efforts. For example, when targeting a highly interconnected group, intervention designers could expect that administering the intervention to a few well-connected individuals will have a strong impact at the community level. In contrast, when targeting a less interconnected group, intervention designers could administer the intervention to more central individuals for a comparable effect. There are, however, two caveats to this example. First, in a lot of cases the most connected individuals are also the most central. Despite this overlap, we argue it is still worthwhile providing more precision to the targeting, since there exist cases in which connected individuals are peripheral, especially in networks with high clustering. Second, we acknowledge that measuring the entire social network might be extremely difficult and costly for intervention designers, especially when it comes to temporally unfolding conversational networks. Surveys to a subset

of individuals in a network in which they could answer questions about how connected they are and their perceived connections among their friends could provide valuable information that could be used for targeting purposes.

Funding

This work was supported by the NSF Science of Learning and Augmented Intelligence Program (grant BCS 2027225).

Author's Note

The data and materials can be found on our open science framework page: <https://osf.io/72ayt/?>.

The data analysis (in python) can be viewed as a jupyter notebook here: <https://github.com/mvlasceanu/CollectiveBeliefs>.

All measures, manipulations, and exclusions are disclosed.

The authors declare no conflicts of interest.

References

- Ajzen, I. (2011). The theory of planned behaviour: Reactions and reflections.
- Bahrami, B., Olsen, K., Bang, D., Roepstorff, A., Rees, G., & Frith, C. (2012). Together, slowly but surely: The role of social interaction and feedback on the build-up of benefit in collective decision-making. *Journal of Experimental Psychology: Human Perception and Performance*, 38(1), 3.
- Benegal, S. D., & Scruggs, L. A. (2018). Correcting misinformation about climate change: The impact of partisanship in an experimental setting. *Climatic change*, 148(1–2), 61–80.
- Borge, M., Ong, Y. S., & Rosé, C. P. (2018). Learning to monitor and regulate collective thinking processes. *International Journal of Computer-Supported Collaborative Learning*, 13(1), 61–92.
- Bouvier, A. (2004). Individual beliefs and collective beliefs in sciences and philosophy: The plural subject and the polyphonic subject accounts: Case studies. *Philosophy of the social sciences*, 34(3), 382–407.
- Cacioppo, J. T., Fowler, J. H., & Christakis, N. A. (2009). Alone in the crowd: the structure and spread of loneliness in a large social network. *Journal of personality and social psychology*, 97(6), 977.

- Coman, A., Momennejad, I., Drach, R. D., & Geana, A. (2016). Mnemonic convergence in social networks: The emergent properties of cognition at a collective level. *Proceedings of the National Academy of Sciences*, 113(29), 8171–8176.
- Cuc, A., Koppel, J., & Hirst, W. (2007). Silence is not golden: A case for socially shared retrieval-induced forgetting. *Psychological Science*, 18(8), 727–733.
- Dovidio, J. F., & Esses, V. M. (2007). Psychological research and public policy: Bridging the gap. *Social Issues and Policy Review*, 1, 5–14.
- Fowler, J. H., & Christakis, N. A. (2010). Cooperative behavior cascades in human social networks. *Proceedings of the National Academy of Sciences*, 107(12), 5334–5338.
- Frankel, R., & Swanson, S. R. (2002). The impact of faculty-student interactions on teaching behavior: An investigation of perceived student encounter orientation, interactive confidence, and interactive practice. *Journal of Education for Business*, 78(2), 85–91.
- Freeman, L. C. (1978). Centrality in social networks conceptual clarification. *Social networks*, 1(3), 215–239.
- Friedkin, N. E., & Johnsen, E. C. (1990). Social influence and opinions. *Journal of Mathematical Sociology*, 15(3–4), 193–206.
- Friedkin, N. E., Proskurnikov, A. V., Tempo, R., & Parsegov, S. E. (2016). Network science on belief system dynamics under logic constraints. *Science*, 354(6310), 321–326.
- Gilbert, M. P. (2000). *Collective belief and scientific change. Sociality and responsibility: New essays in plural subject theory.*
- Hochbaum, G. M. (1958). Public participation in medical screening programs: A socio-psychological study (No. 572). US Department of Health, Education, and Welfare, Public Health Service, Bureau of State Services, Division of Special Health Services, Tuberculosis Program.
- Kim, J. W. (2018). They liked and shared: Effects of social media virality metrics on perceptions of message influence and behavioral intentions. *Computers in Human Behavior*, 84, 153–161.
- King, G., Schneer, B., & White, A. (2017). How the news media activate public expression and influence national agendas. *Science*, 358(6364), 776–780.
- Lewandowsky, S., Ecker, U. K., Seifert, C. M., Schwarz, N., & Cook, J. (2012). Misinformation and its correction: Continued influence and successful debiasing. *Psychological science in the public interest*, 13(3), 106–131.
- Mangels, J. A., Butterfield, B., Lamb, J., Good, C., & Dweck, C. S. (2006). Why do beliefs about intelligence influence learning success? A social cognitive neuroscience model. *Social cognitive and affective neuroscience*, 1(2), 75–86.
- Mesoudi, A. (2011). *Cultural evolution.* University of Chicago Press.
- Milgram, S. (1967). The small world problem. *Psychology Today*, 2, 60–67.
- Momennejad, I., Duker, A., & Coman, A. (2019). Bridge ties bind collective memories. *Nature communications*, 10(1), 1–8.

- Pennycook, G., McPhetres, J., Zhang, Y., Lu, J. G., & Rand, D. G. (2020). Fighting COVID-19 Misinformation on Social Media: Experimental Evidence for a Scalable Accuracy-Nudge Intervention. *Psychological science*, 0956797620939054.
- Richerson, P. J., & Boyd, R. (2005). Not by genes alone: how culture transformed human evolution University of Chicago Press. Chicago, IL, USA.
- Schwitzgebel, E. (2010). Acting contrary to our professed beliefs or the gulf between occurrent judgment and dispositional belief. *Pacific Philosophical Quarterly*, 91(4), 531–553.
- Shariff, A. F., & Rhemtulla, M. (2012). Divergent effects of beliefs in heaven and hell on national crime rates. *PLoS one*, 7(6), e39048.
- Tang, J., Musolesi, M., Mascolo, C., & Latora, V. (2009, August). Temporal distance metrics for social network analysis. In *Proceedings of the 2nd ACM workshop on Online social networks* (pp. 31–36).
- Vlasceanu, M., & Coman, A. (2018). Mnemonic accessibility affects statement believability: The effect of listening to others selectively practicing beliefs. *Cognition*, 180, 238–245.
- Vlasceanu, M., & Coman, A. (2020). Social norms trigger belief change.
- Vlasceanu, M., & Coman, A. (2020). The Effects of Dyadic Conversations on Coronavirus Related Belief Change.
- Vlasceanu, M., Enz, K., & Coman, A. (2018). Cognition in a social context: a social-interactionist approach to emergent phenomena. *Current Directions in Psychological Science*, 27(5), 369–377.
- Vlasceanu, M., Morais, M. J., Duker, A., & Coman, A. (2020). The Synchronization of Collective Beliefs: From Dyadic Interactions to Network Convergence. *Journal of Experimental Psychology: Applied*. Advance online publication. <http://dx.doi.org/10.1037/xap0000265>.
- Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380), 1146–1151.
- Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks. *Nature*, 393(6684), 440.
- Wilkes-Gibbs, D., & Clark, H. H. (1992). Coordinating beliefs in conversation. *Journal of memory and language*, 31(2), 183–194.
- Zhang, J., Le, G., Larochelle, D., Pasick, R., Sawaya, G. F., Sarkar, U., & Centola, D. (2019). Facts or stories? How to use social media for cervical cancer prevention: A multi-method study of the effects of sender type and content type on increased message sharing. *Preventive medicine*, 126, 105751.