The synchronization of collective beliefs:

From dyadic interactions to network convergence

Short title: The synchronization of collective beliefs

Madalina Vlasceanu[ab], Michael J. Morais[bc], Ajua Duker[d], & Alin Coman[a]

[a]Department of Psychology. Princeton University, Princeton, NJ 08544

[b]Princeton Neuroscience Institute, Princeton University, Princeton, NJ 08544

[c]Center for Statistics and Machine Learning, Princeton University, Princeton, NJ 08544

[d]Department of Psychology, Yale University, New Haven, CT 06520

Classification: SOCIAL SCIENCES

Keywords**:** collective beliefs; belief synchronization; social networks; socially-shared retrieval-induced forgetting

Address correspondence to Madalina Vlasceanu, Department of Psychology, Princeton University, Peretsman-Scully Hall, #528, Princeton, NJ 08540, e-mail: mov@princeton.edu; phone: 609-258-7465.

Abstract

Systems of beliefs organized around religion, politics, and health constitute the building blocks of human communities. One central feature of these collectively-held beliefs is their dynamic nature. Here, we study the dynamics of belief endorsement in lab-created 12-member networks using a two-phase communication model. Individuals first evaluate the believability of a set of beliefs, after which, in Phase 1, some networks listen to a public speaker mentioning a subset of the previously evaluated beliefs, while other networks complete a distracter task. In Phase 2, all participants engage in conversations within their network to discuss the initially evaluated beliefs. Believability is then measured both post-conversation and after one week. We find that the public speaker impacts the community's beliefs by altering their mnemonic accessibility. This influence is long-lasting and amplified by subsequent conversations, resulting in community-wide belief synchronization. These findings point to optimal socio-cognitive strategies for combating misinformation in social networks.

Public significance statement

Given that a community's collective beliefs meaningfully impact individual and collective behavior, it is important to ensure that accurate beliefs are disseminated through the population, while scientifically inaccurate beliefs are identified and swiftly corrected. Thus far, attempts to correct misinformation by refuting it have been implemented with limited success. We offer a socio-cognitive framework to reduce believability of inaccurate beliefs in the community by reducing their mnemonic accessibility in a targeted fashion.

Introduction

People's beliefs meaningfully impact their behavior. Religious beliefs about a punishing deity are associated with reduced crime rates (Shariff & Rhemtulla, 2012), beliefs about the flexibility of human abilities cause improvements in academic performance (Mangels et al, 2006), and beliefs about immigration guide voting behavior (Schildkraut, 2010). While an important body of psychological research has explored the relation between beliefs and behavior (Ajzen, 1991), there is scarce research on how communities of individuals synchronize their beliefs. Understanding these synchronization dynamics will reveal not only how to better disseminate accurate beliefs in the population (Osterholm et al, 2015), but will also elucidate how to diminish the spread of misinformation in vulnerable communities (Hough-Telford et al, 2017; Lewandowsky et al, 2012).

Information, and beliefs based on that information, propagate through communities both because of broad exposure to public sources (e.g., politicians, pundits, celebrities), and because individuals within these communities interact with and influence one another. Accordingly, we employ a two-step flow communication model to experimentally investigate the community-wide synchronization of such beliefs (Katz & Lazarsfeld, 1955). Individuals have a set of initial beliefs, then listen to a public speaker that reiterates some of these beliefs, and, finally, communicate with one another within their social networks about these beliefs, (Vlasceanu, Enz, & Coman, 2018). This framework simulates situations in which, for example, beliefs espoused by influential individuals on Twitter or Facebook are then discussed by their followers in subsequent interactions either online or face-to-face (Bhattacharya, Srinivasan, & Polgreen, 2014; Hilbert et al, 2016).

Exposure to social sources of information has been shown to meaningfully impact people's beliefs in surprising ways. Previous research suggested that superficial features of the belief evaluation experience can impact belief endorsement (Gilbert, Krull, & Malone, 1990; Hasher, Goldstein, & Toppino, 1977). For example, information encountered in the past becomes more believable, a phenomenon known as the illusory truth effect (Hasher, Goldstein, & Toppino, 1977; Begg, Anas, & Farinacci, 1992; Ozubko & Fugelsang, 2011; Fazio, Brashier, Payne, & Marsh, 2015). In a typical illusory truth effect paradigm, participants are first asked to assess the truth-status of a series of statements; then, participants are presented with the initial statements again, interspersed with new statements and are asked again to rate the degree to which they think each statement is true. The finding is that repeated statements are judged as more true than novel statements (for a meta-analytic review, see Dechêne, Stahl, Hansen, & Wänke, 2010). The illusory truth effect has been shown to occur due to increased familiarity, through the up-regulation of memory or, in other words, the increase of mnemonic accessibility of information (Ozubko & Fugelsang, 2011). Vlasceanu and Coman (2018) used the same principle to investigate the effect of down-regulating memory, or decreasing mnemonic accessibility, on belief endorsement. They found that mnemonic accessibility influences believability in both directions: a belief that is easier to recall (due to increased mnemonic accessibility) becomes more believable, and a belief that is harder to recall (due to decreased mnemonic accessibility) becomes less believable compared to a baseline belief whose mnemonic accessibility is not manipulated. This research is grounded in a well-established method to alter mnemonic accessibility: retrieval-induced forgetting (RIF; Anderson, Bjork, & Bjork, 1994; Murayama, Mityatsu, Buchli, & Storm, 2014). According to RIF, selectively

remembering previously encoded information results in increased mnemonic accessibility for the remembered information (i.e., rehearsal effect) while at the same time leads to forgetting the information that was unmentioned, but related to the mentioned information. In a typical selective practice paradigm, participants first study category-exemplar pairs (e.g., the "Fruit" category contains the "Apple" and "Pear" exemplars; the "Tree" category contains the "Oak" and "Pine" exemplars) and then selectively practice half of the exemplars from half of the categories in a stem completion task (e.g., "Fruit-Ap___"). Analyses of a final cued-recall test show that practiced items (RP+ items: Fruit-Apple) are remembered better than unpracticed unrelated items (NRP items: exemplars in the "Tree" category)—a rehearsal effect. Unpracticed items related to those practiced (RP− items: Fruit-Pear) are remembered worse than NRP items—a retrieval-induced forgetting effect (RIF). RIF is thought to occur due to inhibitory processes triggered by the response competition during the practice phase (Kuhl, Dudukovic, Kahn, & Wagner, 2007; but see Mensink & Raaijmakers, 1988). Moreover, the selective practice of information occurring in conversational settings has been found to trigger similar effects (Coman, Manier, & Hirst, 2009). Specifically, when a listener monitors a speaker selectively practicing previously encoded information, the listener experiences socially-shared retrieval-induced forgetting (SS-RIF). That is, she forgets information related to what the speaker mentioned in the conversation. This phenomenon occurs because listeners concurrently retrieve the information along with the speaker, which triggers response competition from related memories, just like in the case of RIF (Cuc, Koppel, & Hirst, 2007).

Building on this literature, Vlasceanu and Coman (2018) showed how altering mnemonic accessibility can result in individual-level belief change. They first asked

participants to rate the believability of a set of 24 statements (pretest) organized in 4

categories (e.g., in the "Allergy" category: "Children can outgrow peanut allergies" and

"Some babies are allergic to their mother's milk"; in the "Health" category: "The

majority of people infected with malaria are children" and "Crying helps babies' lungs

develop"). Then, participants listened to a person mentioning 2 statements in each of 2

categories (e.g., "Allergy-Children can outgrow peanut allergies"). Following this

selective practice phase, participants were asked to recall as many of the 24 initial

statements as they could in a cued recall test (i.e., the category was provided). Finally,

participants were asked to rate again (posttest) the believability of the 24 statements.

The memory results showed increased recall rates of selectively practiced statements

(i.e., increased mnemonic accessibility), and decreased recall rates of unpracticed but

related statements (i.e., decreased mnemonic accessibility), both compared to

unpracticed and unrelated statements. Moreover, the belief-level results indicated that

the believability of the practiced statements (e.g., "Allergy-Children can outgrow

peanut allergies") increased from pretest to posttest, and the believability of the

unpracticed but related statements (e.g., "Allergy: Some babies are allergic to their

mother's milk") decreased from pretest to posttest compared to the believability of

unpracticed and unrelated statements (e.g., statements in the "Health" category).

Here, we are interested in whether the findings reported by Vlasceanu and

Coman (2018) occur in a context in which a public speaker addresses large communities

(Phase 1) that will subsequently engage in networked conversations (Phase 2). Previous

research has documented numerous instances in which public speakers significantly

impacted societal level outcomes, from shaping public perceptions of climate change

(Hmielowski et al, 2014), to increasing consumerism (Kumar et al, 2016) or even

influencing voting behaviors (e.g., Oprah Winfrey's public endorsement of Barack Obama was estimated to have led to 1 million additional votes for Obama; Garthwaite & Moore, 2008). Our question is: do public speakers impact communities' beliefs by influencing the mnemonic accessibility of those beliefs? If so, how do conversations that take place after the public speaker's intervention impact people's beliefs? Do people's beliefs synchronize according to the influence of the public speaker and the conversations? Finally, are these effects long lasting? We aim to contribute to the two-step flow communication model by integrating recent psychological advances on the impact of communicative exchanges on memory and belief (Vlasceanu & Coman, 2018; Sperber, 1996).

Methods

Participants. A total of 168 participants (65% women; mean age 21.36 years old) associated with Princeton University were recruited for the study. They participated in the study for either monetary compensation or research credit. Participants were grouped into fourteen 12-member networks. The sample size was determined based on the effect sizes reported in previous research: the average belief suppression effect size was d=0.26, while the average belief rehearsal effect size was d=0.36 (Vlasceanu & Coman, 2018). While a conventional effect size analysis would weigh the false alarm probability against the detection probability and choose a sample size accordingly, we argue that there is an additional consideration involving the network level of analysis. In contrast to the previous research, the current study involves repeated conversational interactions. Because these conversations have been shown to have a cumulative impact on the dependent variable (Coman & Hirst, 2012), and considering the prior reported

effect size (Vlasceanu & Coman, 2018), we would expect an effect size between 0.4 and 0.6. The lower bound, an effect size of 0.4 would require a sample size of 200 participants to be detected. The upper bound, an effect size of 0.6 would only require 42 participants to be detected. Thus, we decided on a stopping rule of 14 networks (168 participants), also taking into consideration previous studies investigating collective-level phenomena that used 14 (Coman et al., 2016) and 12 networks (Momennejad, Duker, & Coman, 2018). This sample size would give us .90 power to detect an effect size of .5 in a between-condition comparison.

Of the 168 participants, 115 (59 in the Experimental and 56 in the Control condition) completed a one-week follow-up survey (69% women; mean age 21.63 years old). The study was approved by the Institutional Review Board at Princeton University.

Stimulus Materials. The materials were borrowed from Vlasceanu & Coman (2018). They consisted of 24 belief statements grouped in four categories (2 myths and 4 correct pieces of information in each category). The myths and the facts were selected based on a pretested dataset collected on MTurk (112 participants), such that they were not statistically significantly different on believability, perceived scientific support, and personal relevance. In addition, the beliefs were correctly categorized as being part of a category by more than 75% of the sample (Vlasceanu & Coman, 2018). The myths were comprised of statements commonly endorsed by individuals as true, but in fact are demonstrably false, whereas the facts were scientifically accurate statements. For example, a myth was that "reading in dim light can damage children's eyes," while a fact was that "children who spend less time outdoors are at greater risk to develop myopia."

Design and Procedure. The 168 participants were split in 14 lab-created communities of 12 participants. Each community was assembled separately, and was comprised by individuals who arrived in the lab at the same time. Sessions were overbooked (by 2-3 participants) to ensure that we had the required community size even if some participants did not show up at the scheduled time. These communities were randomly assigned to either the Control (7 networks) or the Experimental (7 networks) condition. Once assigned to the condition, participants went through 5 experimental phases (Fig.1). In the *study/evaluation phase*, participants were presented with 24 statements that, they were told, are "frequently encountered on the internet." The presentation occurred in a randomized category-blocked fashion and participants were instructed to carefully read these statements. They were also asked to rate the degree to which they believe each statement is accurate, on a scale from 1-(Not at all) to 7-(Very much so) and has scientific support, on a scale from 1-(Definitely not) to 7-(Definitely yes). Next, participants in the Experimental condition listened to an audio of a participant who, supposedly, recalled the information to which s/he was exposed during the experiment in a previous session. In reality, the speaker was a confederate mentioning the statements with minor hesitations to indicate a naturalistic recall. Each participant listened to an audio containing half of the correct statements (i.e., 2 statements) from each of the 4 initially studied categories. Thus, each participant listened to 8 pieces of correct information in the audio, which constituted the RP+ (Retrieval Practice +) items. RP+ items were always correct pieces of information. The RP- (Retrieval Practice -) items were the 16 initially studied items not mentioned in the audio. Participants in the Control condition did not go through a selective practice phase, instead they completed

an unrelated distracter task. Next, participants engaged in sequential dyadic anonymous chat conversations (computer-mediated) as part of the *conversational recall phase*. Each participant took part in a sequence of three or four 3-minute conversations, during which they were asked to collaboratively remember as many of the statements they initially studied as possible: "In this phase you will have a series of chat conversations with other participants about the materials you studied. During these chat conversations you will be asked to jointly remember the information that you studied initially about child rearing. Please be patient and engaged in the task throughout." The conversations were characterized by turn-taking, with virtually all conversational recalls involving collaboration between the interacting partners. A qualitative analysis of the conversations revealed that all of the participants stayed on task throughout the duration of the study. The conversational sequence created a communicative network characterized by 3 clusters, with frequent within-cluster interactions, a network structure that mimics the types of networked interactions one might have in one's community (Watts & Strogatz, 1998) (Fig.1). After the conversational recall phase, participants were randomly presented with the initially read beliefs and were asked to rate them on the same two 7-point scales as before (i.e., degree of belief and scientific support). This rating occurred at two time points, once immediately after the recall phase in the lab (*post-evaluation phase)* and one week later through a Qualtrics link (*follow-up-evaluation phase)*.

All stimuli and procedures were approved by Princeton University's Institutional Review Board (IRB).
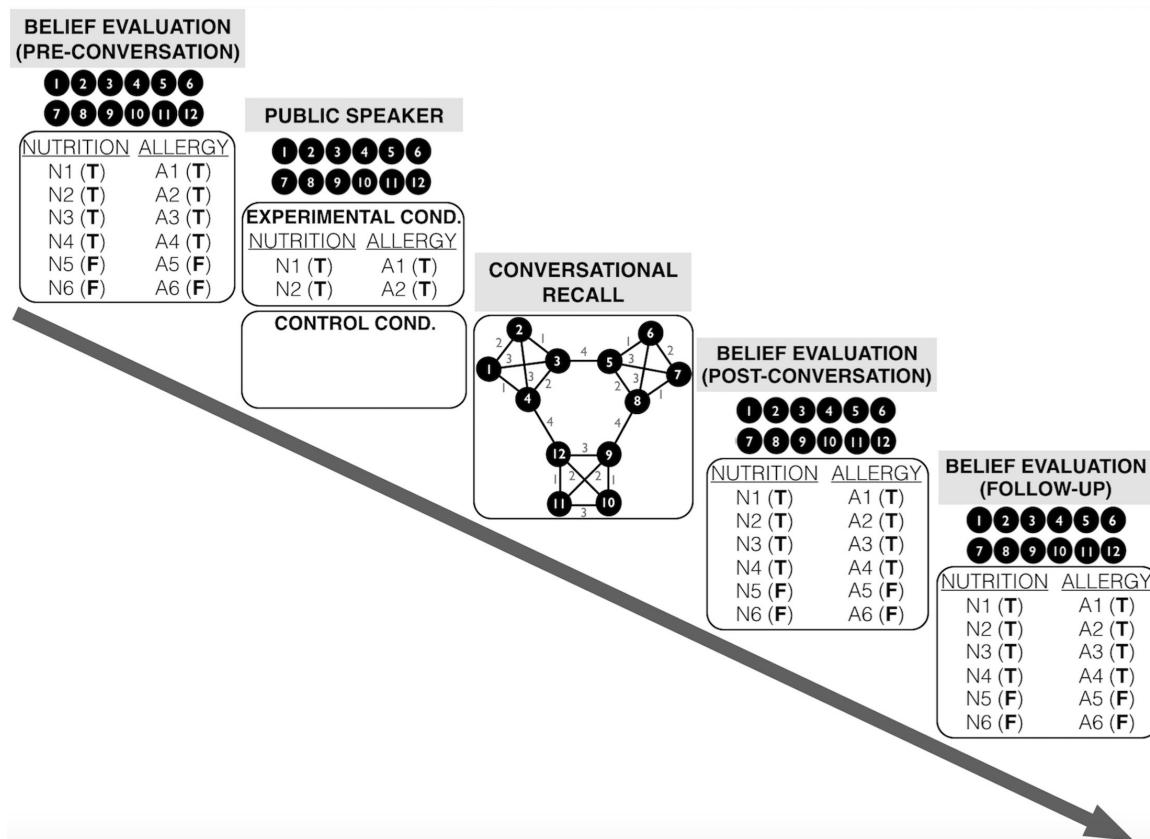
Figure 1. Procedure. In the belief evaluation phase, participants study and evaluate 24 beliefs grouped in 4 categories (two shown here). In each category 4 beliefs are scientifically accurate (T) and 2 are scientifically inaccurate (F). Next, communities are assigned to either the Experimental (public speaker mentions 2 accurate beliefs from each category) or the Control condition (no public speaker). In the conversational recall phase, participants engage in a sequence of dyadic interactions. Circles represent participants and links conversations. Numbers along the links indicate the conversational sequence. Finally, participants are asked to rate the believability of the initially studied beliefs both immediately after the conversational phase and in a one-week follow-up.

We offer three points of clarification about the public speaker. First, we differentiate our usage of the term from the more colloquial usage implying a mass media context. We simply mean that the speaker's message is broadcasted to all participants in the community. Second, the participants were told that the public speaker only selectively remembered some of the beliefs due to time constraints, to diminish the possibility that participants inferred the importance of the beliefs based on whether they were mentioned or not by the public speaker. And third, our main aim in

implying that the speaker in the audio was another participant, was to reduce the impact of source credibility and expertise on our dependent variables.

Analyses and Coding. Each belief was coded as successfully remembered if the recall captured the gist of the original statement. For instance, if for the statement "Eating carrots will make babies' eyesight sharper," participants remembered "Carrots improve vision," their recall was coded as accurate, since it captures the gist of the original statement. Ten percent of the data were double-coded for reliability (Cohen's kappa=0.93) and all disagreements were resolved through discussion between coders. Items that were negated in conversations were excluded from the analyses of both the participant who mentioned the statement as false, and their conversation partner. On average, for each participant, only a small number of beliefs (1.66 out of the 24) were discarded due to negation during the conversational phase, with no difference between the accurate and inaccurate beliefs. We decided to exclude these items because we are interested in observing the impact of mnemonic accessibility on statement believability, which would be contaminated by a discussion of their truth-value.

Supplementary material. The data can be found on the Open Science Foundation website at: https://osf.io/8vjym/

Results

We hypothesized that a belief's mnemonic accessibility would impact its believability. First, we predicted that the public speaker would influence the mnemonic accessibility of the initially studied statements and cause believability changes across the

community. Second, we expected the influence of the public speaker to be amplified in subsequent conversations. Statements mentioned by the public speaker should be more likely to be discussed in ensuing conversations (relative to the Control condition) and would experience an increase in believability from pre to post-conversation (*belief rehearsal effect*). Statements not mentioned by the public speaker, but related to those mentioned should be less likely to be remembered in subsequent conversations (relative to the Control condition), which would result in a decrease in believability from pre to post-conversation (*belief suppression effect*). Finally, these *belief rehearsal and suppression effects* should circumscribe the degree of belief convergence across the community.

The effect of a public speaker on conversational recall and belief endorsement. We first wanted to establish whether there are differences in believability and memorability between the accurate (i.e., facts) and inaccurate beliefs (i.e., myths). Believability was measured by averaging the two highly-correlated evaluations (i.e., perceived accuracy and scientific support, r=0.80) for each belief. We found no differences between the pre-conversational believability scores for myths (M-Myths=4.22, SD=0.79) and for facts (M-Facts=4.19; SD=0.61), *p*=0.74. For memorability, we first computed the recall proportion of each belief by coding the conversational recalls of participants. We then averaged these conversational recall proportions across all rounds of conversation for each participant. This comparison only involved the Control condition, since the recall proportion of the beliefs in the Experimental condition was influenced by the status of the belief (i.e., RP+/RP-). There was no difference between the myths (M-Myths =0.31, SD=0.14) and the facts (M-Facts=0.29, SD=0.13), *p*=0.19. The two types of items were, therefore, indistinguishable from one another, as was found in our preliminary studies

involving these stimulus materials (Vlasceanu & Coman, 2018). Since there was no main effect for a variable coding for truth value of belief nor an interaction with other variables in our analyses, we collapsed across myths and facts here forth.

Next, we investigated whether listening to the public speaker leads to *rehearsal* and *retrieval-induced forgetting effects* on memory. We computed the recall proportion of beliefs mentioned in each participant's conversational recalls. Regardless of who brought up an item in the conversation, it was counted as remembered for both conversational participants, a method typically employed to address the interdependency of individual recalls in conversational interactions (Kashy & Kenny, 1999); this decision was also based on previous studies documenting a similar effect size for speakers as for listeners (Coman, Manier, & Hirst, 2009). For the Experimental condition, we averaged these scores across the rounds that each participant was part of, separating between items mentioned in the audio (RP+ items) and items related to those mentioned in the audio (RP- items). For the Control condition, we separated the beliefs according to the Experimental condition's RP+ beliefs and RP- beliefs, but note that in this condition none of the items were actually practiced before the conversational recalls started. Using these recall proportions as the dependent variable, we conducted a Mixed Factorial ANOVA, with Retrieval Type (RP+ vs. RP-) as a within-subject variable, and Condition (Experimental vs. Control) as a between-subject variable. We found a significant main effect for Retrieval Type, $F_{(1, 160)} = 105.99$, $p < 0.001$, for Condition, $F_{(1, 160)} = 40.26$, $p < 0.001$, and for their interaction, $F_{(1, 160)} = 286.37$, $p < 0.001$. When exploring the interaction, we found that the recall proportion of RP+ items was significantly larger in the Experimental condition (M=0.55, SD=0.18) than in the Control condition (M=0.25, SD=0.14), ($p < 0.001$, Cohen's d=1.86), suggesting that listening to the

public speaker selectively practicing the beliefs leads to increased recall of those beliefs.

Similarly, the recall proportion of RP- items was significantly lower in the Experimental

condition (M=0.27, SD=0.10) than in the Control condition (M=0.32, SD=0.12), ($p <$

0.002, Cohen's d=0.45) indicating that unmentioned beliefs related to those mentioned

were forgotten, relative to the control condition (Fig.2A).
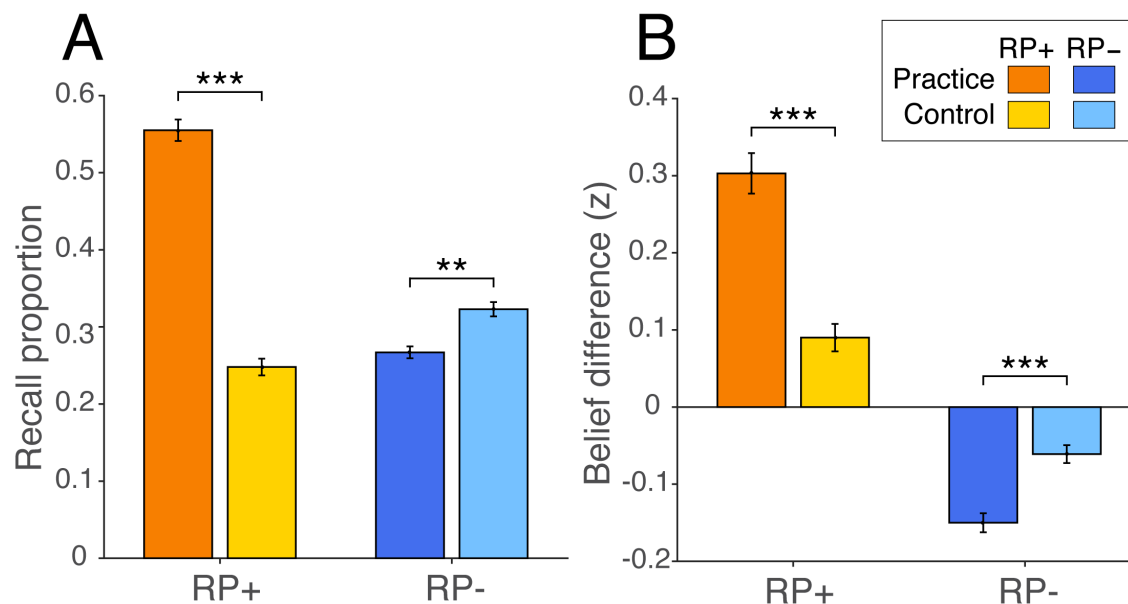
Figure 2. (A) The recall proportion of beliefs averaged over each participant's conversations and separate between the RP+ and RP- beliefs (Experimental) and their corresponding beliefs in the Control condition. **(B)** The post-pre belief difference score, separate between RP+ and RP- beliefs (Experimental) and their corresponding beliefs in the Control condition.

Does this recall pattern result in believability changes? In other words, is

believability dependent on the mnemonic accessibility of the belief? Given the recall

data, we expected a *belief rehearsal effect*, such that practiced beliefs should increase in

believability, and a *belief suppression effect*, such that beliefs related to the practiced ones

should decrease in believability. To explore these predictions, we first standardized the

belief ratings, using z-scores, within-participant. This standardization procedure

allowed controlling for participant-specific particularities of rating scale use, ensuring

any observed effects are not driven by outliers (Fischer & Milfont, 2010). We did so separately for the pre-conversational and post-conversational belief ratings. For each participant, we subtracted the pre-conversational z-score from the post-conversational z-score for each belief, separating between RP+ and RP- beliefs in the Experimental condition and the corresponding beliefs in the Control condition (Fig.3). Using this belief difference score as a dependent variable we conducted a Mixed Factorial ANOVA, with Retrieval Type (RP+ vs. RP-) as a within-subject variable and Condition (Experimental vs. Control) as a between-subject variable. We found a significant main effect for Retrieval Type, $F_{(1, 164)} = 78.22$, $p < 0.001$, for Condition, $F_{(1, 160)} = 31.17$, $p < 0.001$, and for their interaction, $F_{(1, 160)} = 19.31$, $p < 0.001$. Post-hoc analyses showed that the RP+ beliefs became more believable (from pre- to post-conversation) in the Experimental condition (M=0.30, SD=0.34) than in the Control condition (M=0.09, SD=0.23), $p < 0.001$, Cohen's d=0.72. Listening to a public speaker selectively practicing beliefs leads, thus, to increased believability on the part of the listener. Similarly, the RP- items became less believable (from pre- to post-conversation) in the Experimental condition (M=-0.15, SD=0.16) than in the Control condition (M=-0.06, SD=0.15), ($p < 0.001$, Cohen's d=0.58), indicating a *belief suppression effect* for beliefs related to those mentioned (Fig.2B). Listening to a public speaker results, thus, in diminished believability for beliefs related to those mentioned. These results are consistent with the pattern we obtained with an individual-level paradigm (Vlasceanu & Coman, 2018).
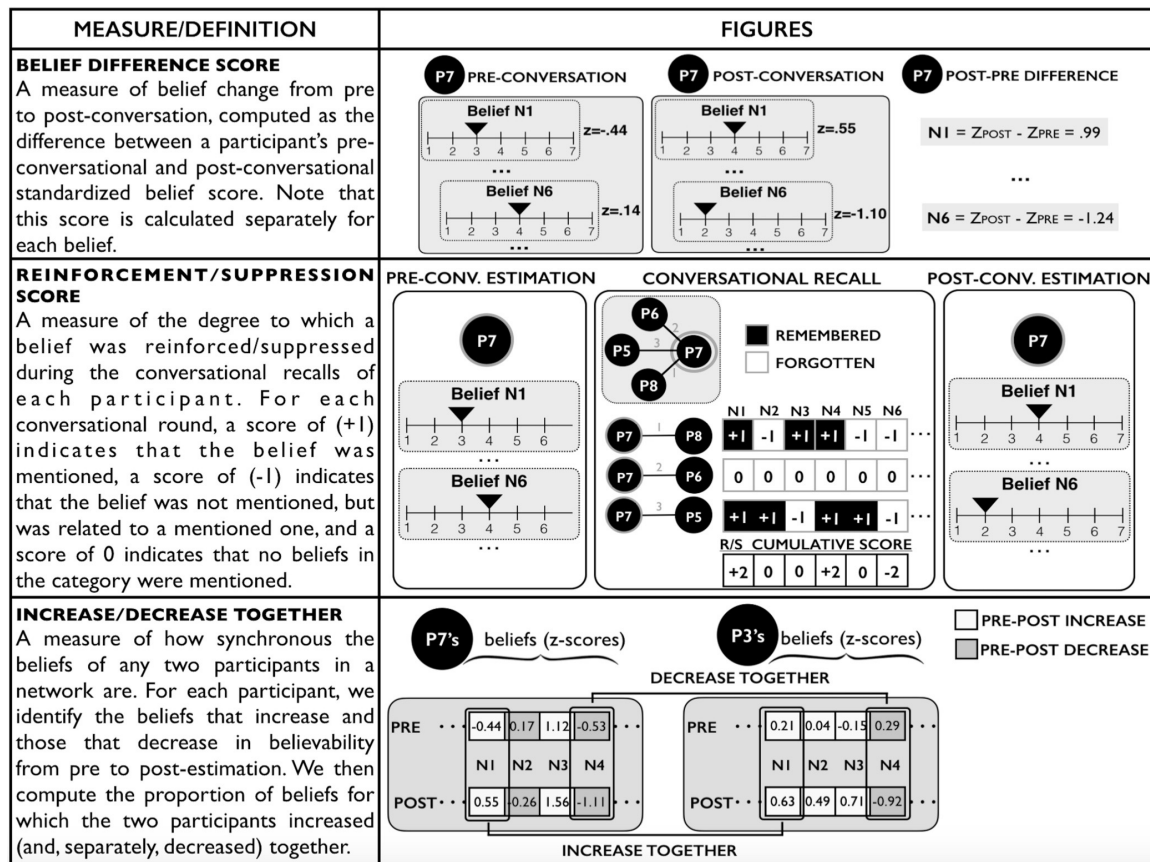
Figure 3. Definitions and illustrative figures for the dependent variables.

The independent and cumulative impact of conversational recall on belief endorsement. We showed that mnemonic accessibility changes triggered by a public speaker's selective practice affects believability. We next assessed the conversational recall's independent impact on believability as well as the cumulative effect of the public speaker and the conversational recall on believability. To test these effects, we analyzed the content of participants' conversations, following previous work (Coman, Momennejad, Drach, & Geana, 2016). We computed cumulative reinforcement/suppression (R/S) scores for each of the 24 initially studied beliefs for each participant as follows. If a belief was mentioned during a conversation, it received a (+1) score on the R/S scale. Similarly, if a belief was not mentioned during a

conversation, but was related to a belief that was mentioned it received a (-1) score on the R/S scale. Unmentioned and unrelated to the mentioned beliefs received a score of 0 on the R/S scale. The final R/S score for each participant was cumulated across the three/four conversations s/he had in the network and was computed separately for each belief. For instance, if a belief was mentioned in all three conversations that a participant had in the network, then it's cumulative R/S score was (+3), while if the belief was part of the category mentioned during all the conversations that the participant was engaged in, but was itself never mentioned in any of the three conversations, then it's R/S cumulative score was (-3). We did not account for the source of the information during the conversation (i.e., who was the speaker and who was the listener) since previous research showed that during conversational recall the speakers and listeners experience similar degrees of rehearsal and retrieval-induced forgetting effects (Coman & Hirst, 2012). We predicted that items with positive cumulative R/S scores will experience a *belief rehearsal effect* such that they will become more believable post-conversation (relative to pre-conversation), while items with negative R/S scores will experience a *belief suppression effect,* such that they will become less believable post-conversation (relative to pre-conversation). To capture these effects, for each belief we computed a belief difference score by subtracting its pre-conversational z-score from its post-conversational z-score (Fig.3). Since there was no participant who had belief z-score values for all nine R/S levels (-4 to 4), we collapsed all the beliefs that had positive R/S scores by averaging the belief z-scores across the R/S scores that ranged from (+1) to (+4). Similarly, we collapsed all negative R/S scores by averaging the belief z-scores across the (-4) to (-1) R/S scores. A positive value for

this belief difference score indicates a *belief rehearsal effect*, whereas a negative value indicates a *belief suppression effect*.

We wanted to investigate whether the conversations had an independent effect on believability. Using the belief difference score as a dependent variable, we ran a Mixed Factorial ANOVA with R/S Item Type as a within-subject variable (Negative R/S; Zero; Positive R/S) and Condition (Experimental vs. Control) as a between-subject variable. We found a main effect for R/S Item Type, $F_{(2, 114)}= 6.10$, $p < 0.003$, but not for Condition ($p = 0.43$). As predicted, the interaction between R/S Item Type and Condition was significant, $F_{(2, 114)}= 3.42$, $p < 0.036$. Post hoc analyses revealed evidence for the *belief suppression effect*, with Negative R/S scores being remembered worse in the Experimental condition (M=-0.12, SD=0.22) than in the Control condition (M=-0.01, SD=0.23), $t(164)=3.14$, $p<0.02$, Cohen's d=0.49. No *belief rehearsal effect* was found, even though the direction of the difference was consistent with the hypothesis that Positive R/S scores would be remembered better in the Experimental condition (M=0.15, SD=0.39) than in the Control condition (M=0.06, SD=0.46) ($p=0.19$, Cohen's d=0.21) (Fig.4A).

This analysis did not differentiate, however, between items mentioned (vs. those related to those mentioned) by the public speaker. If the conversational recall amplified the effects triggered by the public speaker, we should observe higher belief difference scores between the Experimental and Control conditions for RP+ items that had positive R/S scores, and significantly lower belief difference scores between the Experimental and Control conditions for RP- items that had negative R/S scores. And indeed, both differences were statistically significant: $t(143)=2.31$, $p<0.022$, Cohen's d=0.40, and $t(164)=3.02$, $p<0.003$, Cohen's d=0.46, respectively (Fig.4B, 4C). This

indicates that the impact of a public speaker on people's beliefs is stronger if its

influence is further propagated in their subsequent conversations. It is the conjunction

of a public speaker's interventions and people's ensuing conversations that facilitate

both the *belief rehearsal* and *belief suppression effects*.
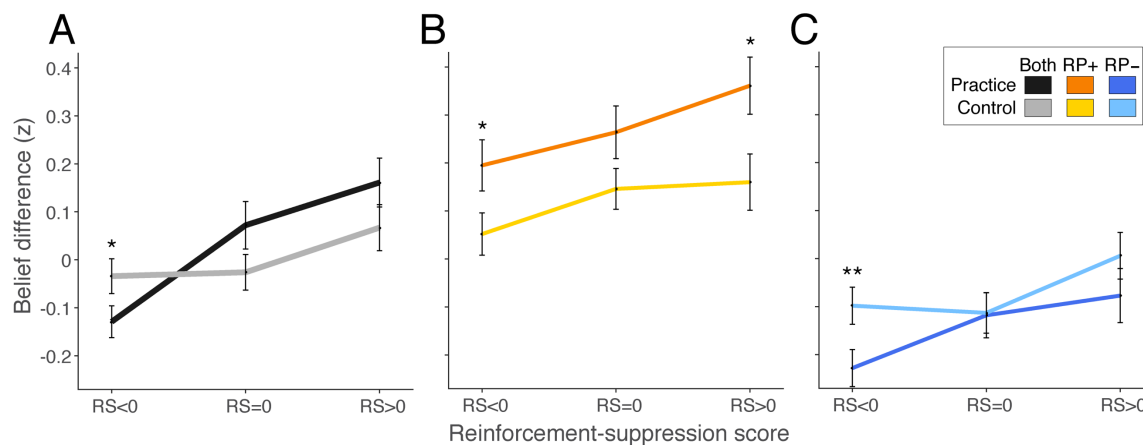
Figure 4. (A) The effect of conversational recall on the post-pre belief difference score. R/S indicates the Cumulative Reinforcement/Suppression score; beliefs with RS<0 were unmentioned in the participant's conversations, but were related to those mentioned, while beliefs with RS>0 were mentioned during the participant's conversations. **(B)** The effect is separated for the beliefs mentioned by the public speaker and their corresponding beliefs in the Control condition. **(C)** The effect is separated for the beliefs that were related to those mentioned by the public speaker and their corresponding beliefs in the Control condition.

Belief synchronization is affected by selective practice. Does the public speaker lead

to belief synchronization at a community level? To test for belief synchronization, we

reasoned that the beliefs mentioned by the speaker (RP+) would be more likely to

increase in believability in unison among community members, while beliefs related to

those mentioned (RP-) would decrease in believability, again, in unison across the

community, relative to the Control condition. We used the post-pre belief difference

scores to compute the proportion of beliefs that increased together (in believability) and

the proportion of beliefs that decreased together between every pair of participants in a

network (Fig.3). We then separated between RP+ and RP- beliefs. For the Control

condition, we computed the same proportions, for the items corresponding to the

RP+/RP- beliefs. Using these pair-wise proportion scores as the dependent variable, we conducted a Mixed Factorial ANOVA with Retrieval Type (RP+ vs. RP-) and Pre-Post Dynamic (Increase vs. Decrease) as within-subject factors and Condition (Experimental vs. Control) as a between-subject variable. We found a significant three-way interaction, $F(1, 900) = 151.39$, $p < 0.001$. As predicted, the proportion of RP+ beliefs that increased together was higher in the Experimental (M=0.37, SD=0.19) than in the Control condition (M=.28, SD=.17), $p<0.001$, Cohen's d=0.50. Conversely, the proportion of RP+ beliefs that decreased together was lower in the Experimental (M=0.15, SD=0.13) than in the Control condition (M=0.24, SD=0.17), $p<0.001$, Cohen's d=0.59. The RP- beliefs exhibited the opposite pattern, such that the proportion of RP- items that decreased together was larger in the Experimental (M=0.39, SD=0.15) than in the Control condition (M=0.32, SD=0.16), $p<0.001$, Cohen's d=0.45, and the proportion of RP- items that increased together was lower in the Experimental (M=0.15, SD=0.11) than in the Control condition (M=0.20, SD=0.14), $p<0.001$, Cohen's d=0.40 (Fig.5). This pattern suggests that exposure to a public speaker triggers community-wide belief synchronization.
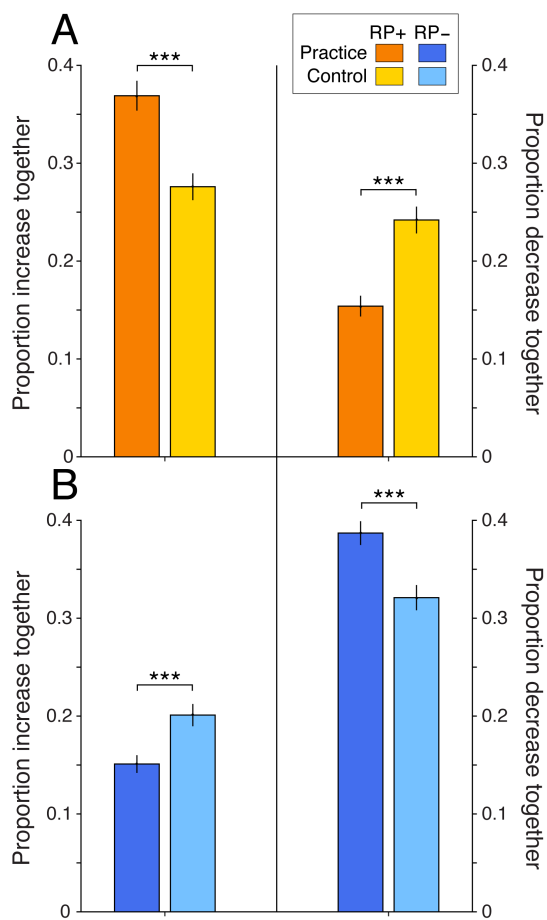
Figure 5. The proportion of beliefs that increase together or decrease together, from pre to post-conversation, averaged across all pairs of participants in the community, and separate for (**A**) RP+ beliefs (Experimental) and their corresponding beliefs in the Control condition and **(B)** RP- beliefs (Experimental) and their corresponding beliefs in the Control condition.

Do these effects endure over time? In summary, we found that a public speaker affects both what people remember and what they believe. Moreover, this impact leads to the synchronization of people's beliefs across the community. But how long-lasting are these effects? Previous research has found that the rehearsal and retrieval-induced forgetting effects can last for at least a week (Storm, Bjork, & Bjork, 2012). No research to date has investigated how long-lasting the *belief rehearsal* and *belief suppression effects* are. In order to answer this question, we conducted a follow-up, 7.31 days (SD=1.29) after

the initial participation in the study. After contacting all 168 participants from the first phase of the study, we collected belief evaluation data from 115 participants (59 participants in the Experimental and 56 participants in the Control condition).

The effect of the public speaker on belief endorsement at follow-up. We first subtracted the pre-conversational belief z-score from its follow-up z-score, separating between RP+ and RP- beliefs in the Experimental condition and the corresponding beliefs in the Control condition. Using this belief difference score as a dependent variable, we conducted a Mixed Factorial ANOVA, with Retrieval Type (RP+ vs. RP-) as a within-subject variable, and Condition (Experimental vs. Control) as a between-subject variable. We found a significant main effect for Retrieval Type, $F_{(1, 113)} = 62.33$, $p < 0.001$, for Condition, $F_{(1, 113)} = 13.52$, $p < 0.001$, and for their interaction, $F_{(1, 113)} = 6.35$, $p < 0.013$. Post-hoc analyses that the RP+ beliefs became significantly more believable (from pre-conversation to follow-up) in the Experimental condition (M=0.29, SD=0.31) than in the Control condition (M=0.14, SD=0.27), $P < 0.005$, Cohen's d=0.52, suggesting that listening to the public speaker selectively practicing beliefs leads to increased believability on the part of the listener one week after exposure. Similarly, the RP- items became marginally less believable (from pre-conversation to follow-up) in the Experimental condition (M=-0.15, SD=0.16) than in the Control condition (M=-0.08, SD=0.18), $P = 0.066$, Cohen's d=0.41, indicating a marginal *belief suppression effect* one week after initial exposure for beliefs related to those mentioned (Fig.6A).
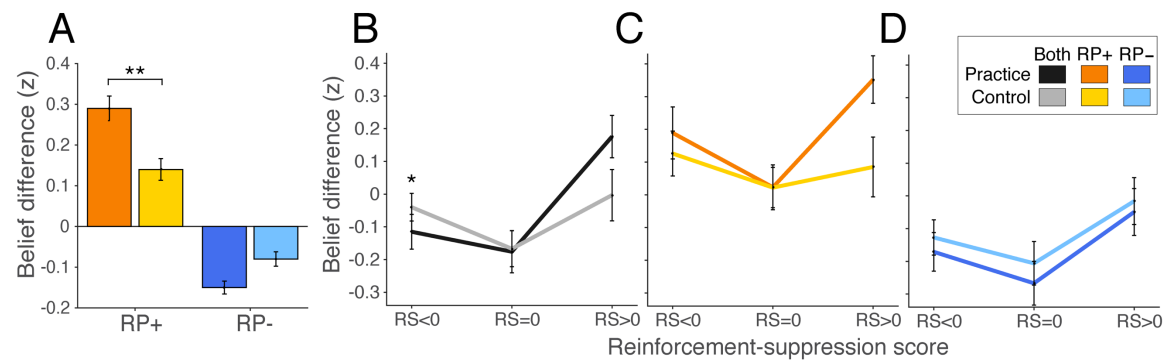
Figure 6 (A) The follow-up/pre belief difference score, separate between RP+ and RP- beliefs (Experimental) and their corresponding beliefs in the Control condition. **(B)** The effect of conversational recall on the follow-up/pre belief difference score. **(C)** The effect is separated for the beliefs mentioned by the public speaker and their corresponding beliefs in the Control condition. **(D)** The effect is separated for the beliefs that were related to those mentioned by the public speaker and their corresponding beliefs in the Control condition.

The independent and cumulative impact of conversational recall on belief endorsement at follow-up. As in previous analyses, in order to explore whether the conversations had an impact on believability at follow-up, we employed a Mixed Factorial ANOVA with R/S Item Type as a within-subject variable (Negative R/S; 0; Positive R/S) and Condition (Experimental vs. Control) as a between-subject variable. The dependent variable was the difference between a belief's z-score at follow-up and its z-score in the pre-conversational phase. We found a main effect for R/S Item Type, $F_{(2, 77)}$= 5.46, $p < 0.006$, but not for Condition ($p = 0.67$). As predicted, we found an interaction between R/S Item Type and Condition, $F_{(2, 114)}$= 3.72, $p < 0.029$. Posthoc tests revealed that the *belief suppression effect* was significant, with Negative R/S scores becoming less believable in the Experimental condition (M=-0.11, SD=0.24) than in the Control condition (M=0.00, SD=0.21), t(113)=2.85, p<0.005, Cohen's d=0.49. We also found a *belief rehearsal effect*, with Positive R/S scores becoming more believable in the Experimental condition (M=0.14, SD=0.33) than in the Control condition (M=-0.03, SD=0.47) t(113)=2.20, *p*<0.03,

Cohen's d=0.42. This indicates that people's conversations meaningfully impact the believability scores of the studied beliefs one week after they occur (Fig.6B).

We showed that the effect of the public speaker on people's beliefs is amplified in their subsequent conversations as measured immediately after the public speaker's intervention. Is this amplification effect long-lasting? If the conversational recall facilitated the effects triggered by the public speaker, we should observe significantly higher belief difference scores between the Experimental and Control conditions for RP+ items with positive R/S scores, and significantly lower belief difference scores between the Experimental and Control conditions for RP- items with negative R/S scores. Neither the *belief rehearsal* (p=0.10), nor the *belief suppression* (p=0.20) effects reached statistical significance at follow-up. These results suggest that the cumulative impact of the public speaker and ensuing conversations on beliefs loses strength with time (Fig.6C/6D).

Is belief synchronization long-lasting? As in previous analyses, we used the increase together/decrease together pair-wise proportion scores as a dependent variable. This time, the increase/decrease was computed using the change between the belief's pre-conversational z-scores and the follow-up z-scores. We conducted a Mixed Factorial ANOVA with Retrieval Type (RP+ vs. RP-) and Pre-Post Dynamic (Increase vs. Decrease) as within-subject variables and Condition (Experimental vs. Control) as a between-subject variable. The three-way interaction was significant, $F(1, 309) = 29.11$, $p < 0.001$. Post-hoc analyses revealed that the synchronization pattern observed in the post-conversational analysis was also present at follow-up. The proportion of RP+ beliefs that increased together was higher in the Experimental condition (M=0.38,

SD=0.18) than in the Control condition (M=0.29, SD=0.17), $p<0.001$, Cohen's d=0.51,

while the proportion of RP+ beliefs that decreased together was lower in the

Experimental (M=0.16, SD=0.14) than in the Control (M=0.19, SD=0.15) condition,

$p<0.04$, Cohen's d=0.21. The proportion of RP- items that increased together was lower

in the Experimental (M=0.17, SD=0.10) than in the Control condition (M=0.21, SD=0.13),

$p<0.001$, Cohen's d=0.34, while the proportion of RP- items that decreased together was

marginally higher in the Experimental condition (M=0.34, SD=0.13) than in the Control

condition (M=0.32, SD=0.16), $p=0.083$, Cohen's d=0.14 (Fig.7).
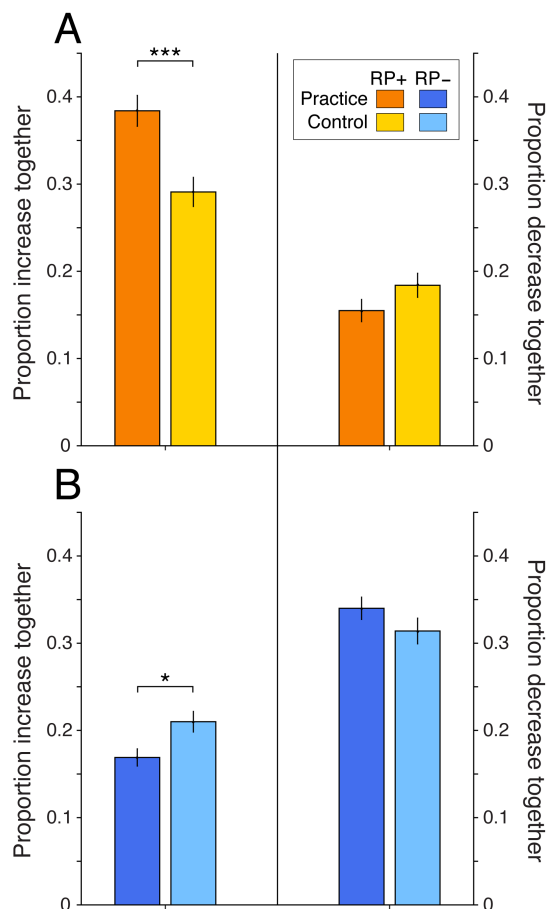
Figure 7. The proportion of beliefs that increase together or decrease together, from pre to follow-up, averaged across all pairs of participants in the community, and separate for the (**A**) RP+ beliefs (Experimental) and their corresponding beliefs in the Control condition and (**B**) RP- beliefs (Experimental) and their corresponding beliefs in the Control condition.

Discussion

We have shown that a public speaker influences a community's beliefs by impacting the mnemonic accessibility of those beliefs. Beliefs mentioned by the public speaker become more believable, while beliefs related to those mentioned become less believable compared to a situation in which no public speaker existed. Importantly, these effects are amplified by conversations within networks, revealing a cumulative impact of the public speaker and the conversations. These effects also regulate the degree of belief convergence across communities, with networks exposed to a public speaker synchronizing their beliefs more than control networks. Finally, we observe the effects, with sizes typical of similar paradigms (Murayama et al., 2014; Vlasceanu & Coman, 2018; Coman et al., 2016), lasting for at least one week.

The strength of our approach to investigating collective belief endorsement is its controlled, experimental nature, which adds meaningfully to research investigating social network data from platforms such as Twitter, which is mostly correlational (Hilbert et al, 2016). It is important to note, however, that in real-world situations the social context in which a public speaker communicates information to an audience is far more complex than tested in the current investigation. Factors such as expertise, credibility, and similarity of the public speaker would likely affect the degree to which the audience integrates or resists the information that is conveyed (Fiske & Taylor, 2015). Moreover, real-world conversations can also be more complex than the ones elicited in our experiment, especially if they contain strong opinions and evaluations. However, our goal was to provide empirical evidence of the impact of mnemonic

accessibility on collective beliefs under minimal conditions, which ensure a highly controlled experimental design. The investigation of these minimal conditions allows us to claim that it is mnemonic accessibility, rather than the characteristics of the social source of information, that triggered the observed belief change. Now that we established this effect, future studies can investigate the impact of variables such as the public speaker's expertise, credibility, and similarity on collective belief endorsement.

Remarkably, a 20-minute session in which participants were exposed to a public speaker and then engaged in networked conversations was sufficient to trigger belief change that lasted for at least one week. The two sources of social influence (i.e., public speaker and conversational interactions) had an independent impact on people's beliefs, both immediately post-conversation and at follow-up. The influence of these sources was also cumulative, such that beliefs mentioned by the public speaker that were also rehearsed in the conversations became most believable among all beliefs, while beliefs related to those mentioned by the public speaker, which were also related to beliefs discussed in people's conversations became least believable. This cumulative effect was found to be temporally limited, though diminishing at follow-up, which points to the boundary conditions of the impact of mnemonic accessibility on believability. For the cumulative effect of the two sources of influence to be long-lasting, additional factors might need to be implemented. Repeated interactions over time following the public speaker's intervention (Centola, 2010) and the ideological consistency between the participants' beliefs and those espoused by the public speaker (Coman & Hirst, 2012) constitute two such factors that we plan to investigate in future research.

In this experiment, we created a conversational network structure that mimicked the main characteristics of real-world social networks: clustered communities

characterized by frequent within-cluster interactions (Watts & Strogatz, 1998). This methodology allowed us to situate individual-level cognitive processes in a framework (Vlasceanu, Enz, & Coman, 2018) aimed at programmatically investigating how micro-level cognitive processes (i.e., rehearsal, retrieval-induced forgetting) could give rise to collective-level large-scale phenomena (i.e., the synchronization of beliefs across communities). We have not undertaken, however, a programmatic investigation of how the network structure, or temporal sequencing of conversations impacts collective-level outcomes. Manipulating the topological (Coman et al, 2016) and temporal (Momennejad, Duker, & Coman, 2018; Li et al, 2017) features of a community's conversational network could lead to significant advances in understanding how collective beliefs are formed and maintained in networked communities.

Moreover, in the current study participants were exposed to information organized in meaningful categories. While organizing information into structured categories is a naturally occurring process (Chapman, 1967), there may be variation in how individuals spontaneously group information into categories. Thus, a noteworthy expansion of this investigation could be exploring how these idiosyncrasies interact with processes we explored herein.

Finally, our findings provide promising possibilities for interventions aimed at countering the spread of inaccurate beliefs in vulnerable communities. A typical intervention targeted at reducing the spread of misinformation involves refuting misinformation by overtly discussing its false nature (Wegner, Wenzlaff, Kerker, & Beattie, 1981). But refutations have been found to reinforce misconceptions (Lewandowsky et al, 2012), especially when individuals are confronted with a complex informational environment (Lamb, King, & Kling, 2003; Contractor & DeChurch, 2014)

and when they are ideologically committed (Nyhan & Reifler, 2010; Flaxman, Goel, & Rao, 2016). We have shown here that in order to diminish the believability of inaccurate beliefs, one does not necessarily need to discuss them. It is sufficient to trigger response competition from the part of inaccurate beliefs by repeatedly broadcasting conceptually related accurate beliefs in the population (Vlasceanu & Coman, 2018). This will result in the suppression of misinformation, which, as we have shown here, will reduce their believability across the entire community.

References

Ajzen, I. (1991). The theory of planned behavior. *Org Behavior and Human Decision Processes*.

    50 (2): 179–211.

Anderson, M. C., Bjork, R. A., & Bjork, E. L. (1994). Remembering Can Cause Forgetting:

    Retrieval Dynamics in Long-Term Memory. *J of Exp Psychol: LMC, 20*(5), 1063–1087.

Begg, I. M., Anas, A., & Farinacci, S. (1992). Dissociation of processes in belief: Source

    recollection, statement familiarity, and the illusion of truth. Journal of Experimental

    Psychology: General, 121(4), 446.

Bhattacharya, S., Srinivasan, P., & Polgreen, P. (2014). Engagement with health agencies on

    twitter. *PLoS One*, *9*(11), e112235.

Carpenter, C.J. (2010). A Meta-Analysis of the Effectiveness of Health Belief Model Variables

    in Predicting Behavior. *Health Comm., 25*:8, 661-669.

Centola D (2010) The spread of behavior in an online social network experiment. *Science,*

    329(5996):1194–1197.

Coman, A., Manier, D., & Hirst, W. (2009). Forgetting the unforgettable through conversation:

    Socially shared retrieval-induced forgetting of September 11 memories. Psychological

    Science, 20(5), 627-633.

Coman, A., Momennejad, I., Drach, R. D., & Geana, A. (2016). Mnemonic convergence in

    social networks: The emergent properties of cognition at a collective level. *PNAS, 113*

    (29), 8171-8176.

Coman A, Hirst W (2012) Cognition through a social network: The propagation of induced

    forgetting and practice effects. *J Exp Psychol: Gen 141*(2):321–336.

Contractor, N. S., & DeChurch, L. A. (2014). Integrating social networks and human social

      motives to achieve social influence at scale. *PNAS, 111*, 13650–13657.

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.

Dechêne, A., Stahl, C., Hansen, J., & Wänke, M. (2010). The truth about the truth: A meta-

      analytic review of the truth effect. Personality and Social Psychology Review, 14(2),

      238-257.

Fazio, L. K., Brashier, N. M., Payne, B. K., & Marsh, E. J. (2015). Knowledge does not protect

      against illusory truth. *J of Exp Psychol: General, 144,* 993-1002.

Fiske, S. T., & Taylor, S. E. (2013). *Social cognition: From brains to culture*. Thousand Oaks.

Flaxman, S.and Goel, S. and Rao, J.M. (2016). Filter Bubbles, Echo Chambers, and Online News

      Consumption. *Public Opinion Quarterly, 80*, 298–320.

Garthwaite, C., & Moore, T. (2008). The role of celebrity endorsements in politics: Oprah,

      Obama, and the 2008 democratic primary. Department of Economics, University of

      Maryland, 1-59.

Gilbert, D. T., Krull, D. S., & Malone, P. S. (1990). Unbelieving the unbelievable: Some

      problems in the rejection of false information. *JPSP, 59*, 601–613.

Hasher, L., Goldstein, D., & Toppino, T. (1977). Frequency and the conference of referential

      validity. *J of Verbal Learning and Verbal Behavior, 16*, 107-112.

Heroux, M. (2017). Cohen's d: How to interpret it? Retrieved from:

      https://scientificallysound.org/2017/07/27/cohens-d-how-interpretation/

Hilbert, M., Vásquez, J, Halpern, D., Valenzuela, S., Arriagada, E. (2016). One Step, Two Step,

   Network Step? Complementary Perspectives on Communication Flows in Twittered

   Citizen Protests. *Social Science Computer Review*, 35 (4), 444 – 461.

Hmielowski, J. D., Feldman, L., Myers, T. A., Leiserowitz, A., & Maibach, E. (2014). An attack

   on science? Media use, trust in scientists, and perceptions of global warming. Public

   Understanding of Science, 23(7), 866-883.

Hough-Telford, C., Kimberlin, D. W., Aban, I., Hitchcock, W. P., Almquist, J., Kratz, R., &

   O'Connor, K. G. (2016). Vaccine delays, refusals, and patient dismissals: a survey of

   pediatricians. *Pediatrics*, e20162127.

Katz & Lazarsfeld (1955). Personal Influence. *New York: Free Press*.

Kashy, D. A., & Kenny, D. A. (1999). The analysis of data from dyads and groups. In H. T. Reis

   & C. M. Judd (Eds.), *Handbook of research methods in social psychology*. New York:

   Cambridge University Press.

Kuhl, B. A., Dudukovic, N. M., Kahn, I., & Wagner, A. D. (2007). Decreased demands on

   cognitive control reveal the neural processing benefits of forgetting. *Nature neuroscience*,

   10(7), 908.

Kumar, A., Bezawada, R., Rishika, R., Janakiraman, R., & Kannan, P. K. (2016). From social to

   sale: The effects of firm-generated content in social media on customer behavior. Journal

   of Marketing, 80(1), 7-25.

Lamb, R., King, J. L., & Kling, R. (2003). Informational environments: Organizational contexts

   of online information use. *Journal of the American Society for Information Science and

   Technology, 54*, 97-114.

Lewandowsky, S., Ullrich K. H., Ecker, Seifert, C.M., Schwarz, N., Cook, J. (2012).

Misinformation and Its Correction: Continued Influence and Successful Debiasing. *Psychol Sci in the Public Interest*, 13(3), pp. 106 – 131.

Li, A., Cornelius, S. P., Liu, Y. Y., Wang, L., & Barabási, A. L. (2017). The fundamental advantages of temporal networks. *Science*, *358*(6366), 1042-1046.

Mangels, J. A., Butterfield, B., Lamb, J., Good, C., Dweck, C. S. (2006). Why do beliefs about intelligence influence learning? A social-cognitive neuroscience model. *SCAN,* 1 (2): 75-86.

Mensink, G. J., & Raaijmakers, J. G. (1988). A model for interference and forgetting. Psychological Review, 95(4), 434.

Momennejad, I., Duker, A., & Coman, A. (2017). The Ties that Bind Networks: Weak Ties Facilitate the Emergence of Collective Memories. *arXiv preprint arXiv:1705.07185*.

Murayama, K., Mityatsu, T., Buchli, D., Storm, B. C. (2014). Forgetting as a consequence of retrieval: a meta-analysis of retrieval-induced forgetting. *Psychol Bull, 140*(5), 1383–1409.

Nyhan, B., & Reifler, J. (2010). When corrections fail: The persistence of political misperceptions. *Political Behavior*, 32(2), 303-330.

Osterholm MT, Moore KA, Kelley NS, Brosseau LM, Wong G, Murphy FA, et al. (2015). Transmission of Ebola viruses: what we know and what we do not know. *MBio. 6.*

Ozubko, J. D., & Fugelsang, J. (2011). Remembering makes evidence compelling: Retrieval from memory can give rise to the illusion of truth. Journal of Experimental Psychology: Learning, Memory, and Cognition, 37(1), 270.

Shariff, A.F., & Rhemtulla, M. (2012). Divergent Effects of Beliefs in Heaven and Hell on National Crime Rates. *PLoS ONE,* 7 (6).

Schildkraut, D.J. (2010). *Americanism in the Twenty-First Century: Public Opinion in the Age of Immigration*. Cambridge University Press.

Sperber, D. (1996). *Explaining culture: A naturalistic approach* (Vol. 323). Oxford: Blackwell.

Storm, B. C., Bjork, E. L., & Bjork, R. A. (2012). On the durability of retrieval-induced forgetting. Journal of Cognitive Psychology, 24, 617-629.

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., 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.

Watts, D. J.; Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks. *Nature*. 393 (6684): 440–442.

Wegner, D. M., Wenzlaff, R., Kerker, R. M., & Beattie, A. E. (1981). Incrimination through innuendo: Can media questions become public answers? *JPSP*, *40*(5), 822.