Exchanging and updating opinions and beliefs reinforces echo chambers

Hiro Kataoka[*1*2] Jérôme Euzenat[*2] Koji Hasebe[*1]

[*1] Department of Computer Science, University of Tsukuba, Tsukuba, Japan
[*2] Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG, F-38000 Grenoble, France

Echo chamber, the state in which agents are split into groups sharing the same opinions, is a well-known phenomenon in social networks. Opinion dynamics models have been proposed to explain how the phenomenon occurs through agents revising their opinions. However, social network users also exchange beliefs supporting their opinions. This has not been taken into account. In this paper, we extend an existing opinion dynamics model by allowing agents to exchange and update both beliefs and opinions. The process of updating beliefs is described based on the classical belief revision theory. Beliefs and opinions can influence each other guided by values that agents share. We compare opinion propagation with respect to belief influence. Simulation results show that echo chambers are reinforced by interactions between opinions and beliefs.

1. Introduction

Echo chamber is one of the well-known phenomena in social networks. It refers to a state in which people interact primarily with those who have similar opinions [Mahmoudi 24], which lead them to believe that their opinions are universally valid. Network segregation, which means that the network is splited into less connected communities, and opinion polarization, which means that agents in the same community share the same opinions. Opinion dynamics models have been applied to explain how echo chambers are formed by aligning agents' opinions with those that they receive from others. For example, it is shown that echo chambers raise through opinion updates with ignoring different opinions and modification of the network [Sasahara 20], similarly to existing social networks.

However, this model does not take into account the effect of what agents believe. Hereafter, opinions and beliefs are different: opinions of an agent denote how positive they are toward corresponding topics while beliefs represent what it considers as true. People exchange and revise not only their opinions but also their beliefs through interactions. For example, beliefs can be communicated to justify opinions: "I have an opinion O toward ϕ because of my beliefs B".

Some frameworks can be applied to express agents' beliefs. For example, how to formalize the way beliefs should be merged so that it satisfies some ideal characteristics is studied in the field of belief revision [Schwind 15].

Based on the these observations, we introduce a model defining interactions between opinions and beliefs. In this model, opinions and belief are propagated through existing models [Sasahara 20, Schwind 15] and they interact together through a mechanism aiming at ensuring their coherence. Using the proposed model, we aim at showing whether allowing agents to interact with others based on beliefs leads to more polarized opinions and more segregated network or not.

Multi-agent simulations of this model show the emergence of echo chambers when agents interact based solely on their beliefs or their opinions. Moreover, agents' final opinions and beliefs are more polarized when opinions and beliefs are allowed to interact.

The remainder of this paper is organized as follows. Section 2 shows how agents behavior in the proposed model is described in an intuitive scenario. Section 3 provides some existing and relevant models to ours. In Section 4, we show hypotheses we will verify in this paper. The proposed model is defined in Section 5. Then, the experimental settings are described in Section 6. Section 7 presents the results from the experiments.

2. Intuitive Scenario

Before we introduce our model, we show an informal scenario that describes how agents exchange their opinions and beliefs in our model and how they form echo chambers through the interactions.

Alice, Bob and Charles are discussing whether the city they are living in should construct a tram network or not. Alice is for the construction plan because she believes that this would reduce traffic jams. On the other hand, Bob is against the plan because he believes that this will lead to raise taxes. Charles is strongly against the plan because he believes that trams are not green transportation.

Their discussion starts from listening each other's opinion and beliefs. Whether a person take into account what another explains or not depends on how close their opinions and beliefs are. Here, assume that Bob and Alice have close enough opinions and beliefs, while Alice ignores Charles because theirs are too different. Then, Alice listens to what Bob explains and updates her opinion (i.e., how positive she considers the topic) and beliefs (i.e., what she believes). For example, she may become slightly more negative toward the topic. Also, she may believe that if their city constructs the tram network then there will be less traffic jams and higher taxes.

If she considers higher tax as bad based on her values, or at least worse than traffic jams, she will be more negative about the topic. On the other hand, she may change her

Contact: Hiro Kataoka, University of Tsukuba, kataoka@mas.cs.tsukuba.ac.jp

beliefs about how much traffic jams will be reduced to better support her new opinion.

Finally, Alice can decide to stop listening to Charles and starts listening to someone else (David for example) instead. If she decides to do so, when she updates her opinion and beliefs next time, she will take Bob's and David's opinion and beliefs into consideration.

If they continue interacting in this way, the four people can be splited in two groups such that people listen to other people in the same group but ignore someone in different groups. People in the same group will share the same (or enough similar) opinions and beliefs while what are shared in different groups will be different. Such a state is called as "echo chamber".

Considering the story above, some questions can be asked: if a large number of agents behave as described above, what can be observed in the end? Agents who update their opinions from the opinions of other agents which are close to theirs will form echo chambers [Sasahara 20]. Does introducing beliefs help fighting against echo chambers or does it reinforce them? Or, does it have no relation with echo chambers and thus existing opinion dynamics models are enough to capture the phenomenon?

3. State of the art

The analysis of such scenarios have been considered from the standpoint of opinion dynamics and belief revision. Thus, related models in these areas are introduced in this section.

3.1 Opinion dynamics models and echo chambers

Echo chambers in social networks describe a situation in which each person interacts only with someone who have opinions similar to theirs [Mahmoudi 24]. It is observed in many online platforms [Cossard 20, Gao 23]. One of the classical models that can be applied to it is the bounded confidence models [Deffuant 00, Hegselmann 02]. Let $O_a^t \in \mathbb{R}$ be the opinion of the agent a at time t (t is omitted if it is clear). In [Deffuant 00], two agents a and a' revise their opinions to get new ones if the difference between their opinions, i.e., $|O_a^t - O_{a'}^t|$, is equal to or less than a threshold $\varepsilon \in \mathbb{R}$. The updating rule can be described as:

$$O_a^{t+1} = O_a^t + \mu(O_{a'}^t - O_a^t),$$
$$O_{a'}^{t+1} = O_{a'}^t + \mu(O_a^t - O_{a'}^t),$$

such that μ is a constant between 0 and 0.5 [Dong 18]. Based on this model, it is showed that low tolerance for dissimilar opinions prevents agents' opinions from converging.

Nowadays many opinion dynamics models have been proposed to explain how echo chambers are formed [Springsteen 24, Cinus 22, Ohagi 24]. A model with opinion threshold, timelines, and network rewriting has been proposed to model online social networks, and particularly Facebook, realistically [Sasahara 20]. In this model, each agent follows at least an other agent. Let A be the set agents and $N_a \subseteq A$ be the set of agents followed by agent $a \in A$. At each iteration, one agent (a) is chosen randomly from A and performs the following:

1. retrieve those posts from a's timeline which opinion is within the threshold (ε) from a's own opinion;
2. update a's opinions with opinions selected at step 1;
3. unfollow an agent $a' \in N_a$ who sent a discordant post and follow another agent $a'' \in A \setminus (N_a \cup \{a\})$ if possible;
4. submit a new post or repost an existing one in a's timeline.

The steps 3 and 4 are executed probabilistically. It is observed that the smaller the threshold, the more echo chambers appear. Moreover, it is reported that unfollowing and following accelerate their appearance.

Nagura and Akiyama [Nagura 25] show that the fewer topics that are discussed by agents and the more important they are, the more echo chambers they form using an extension of the model proposed by [Sasahara 20].

3.2 Belief change

In the field of belief change, many operators describe how to revise beliefs with others have been proposed. The belief revision games (BRG) [Schwind 15] define some procedures, similar to those of opinion dynamics, for merging. One of the example of operators R^{*1} defined in BRG merges n beliefs $B_1, B_2, \ldots, B_n$ into new beliefs. It is defined as:

$$\mathcal{M}(R(B_1, B_2, \ldots, B_n)) = \operatorname*{argmin}_{m \in \mathcal{M}(\top)} \sum_{k=1}^{n} d(m, B_k) \quad (1)$$

such that $\mathcal{M}(\phi)$ is the set of models of ϕ and $d(m, \phi)$ is the minimum number of assignment of true and false to the atomic propositions which should be changed from it in m to make ϕ as true. Because the postulate **(IC3)** [Konieczny 02] is satisfied by R, the merging result is syntax-independent [Schwind 15]. Moreover, R returns consistent beliefs [Schwind 15]. In BRG, some properties including stability of beliefs are showed theoretically.

BeliefFlow [Schwind 24] is another framework of belief revision. It enables beliefs to be updated asynchronously. In this framework, it is proven that agents reach a consensus if the network they form is strongly connected.

The network knowledge base [Gallo 21] also considers propagating knowledge through the exchange of literals. One of its specific feature is the ability that agents have to modify the agent network.

3.3 Models evolving opinions and beliefs

Very few models can be viewed as connecting opinions and beliefs. For example, in belief systems with logic constraints [Friedkin 16, Nedić 19], beliefs are represented as real values and updated using influence matrices between statements which is the same updating rule as [DeGroot 74]. Thus, beliefs in this context can be viewed as opinions. Also, [Botan 19] uses opinions which are represented as boolean and updated using integrity constraints

*1 This operator is equivalent to $R^3_{\Delta^{d_H},\Sigma}$ in [Schwind 15]. In this paper, this operator is expressed as R to make the notation simpler because no other operators from BRG will be used.

which are usually discussed as logical formulae. Thus, opinions in this context can be viewed as beliefs.

4. Hypotheses

To answer the question that is raised at the end of the Section 2, we have three hypotheses.

First, the proposed model is defined based on the model in [Sasahara 20]. Because of this, it is expected that the model can reproduce their results, if opinions are updated independently from beliefs. Thus, the first hypothesis is:

H0 If opinions evolve independently from beliefs, the less strict agents are in selecting those to pay attention to from opinions, i.e., the larger ε is, the closer the final opinions are and the less opinions-based communities agents form.

Next, it is expected that, using a similar procedure, the same will occur to beliefs. Thus, the second hypothesis is:

H1 If beliefs evolve independently from opinions, the less strict agents are in selecting those to pay attention to from beliefs, i.e., the larger δ is, the closer the final beliefs are and the less beliefs-based communities agents form.

Finally, it is expected that agents will be ignored more easily based on both of opinions and beliefs than based on either opinions or beliefs independently. This can strengthen the extent of echo chambers. Thus, the final hypothesis is:

H2 Connecting opinions and beliefs reinforce the polarization and segregation effects of echo chambers.

5. Theoretical framework: a model for opinion and belief propagation and integration

We propose a model in which opinions and beliefs are distinguished: opinions represent how positive an agent considers a topic, while beliefs describe what it regards as true.

5.1 Preliminaries

Let A be the set of agents. Each agent $a \in A$ has its opinion O_a^t toward a topic ϕ and beliefs B_a^t at time t. Opinions are represented as numbers between 0 and 1, i.e., $O_a^t \in \mathcal{O} = [0,1] \subseteq \mathbb{R}$ is the opinion space. Similarly, beliefs are represented as propositional formulae, i.e., $B_a^t \in \mathcal{L}_\mathcal{P}$ such that $\mathcal{L}_\mathcal{P}$ is the set of propositional formulae constructed from the set $\mathcal{P}$ of atoms. The topic ϕ is also represented as a propositional formula, i.e., $\phi \in \mathcal{L}_\mathcal{P}$. Note that we consider only one topic in this paper.

Each agent $a \in A$ has a non empty set of neighbors $N_a^t \subseteq A \setminus \{a\}$ at time t. The directed network agents form at time t can be denoted as $G^t = (A, N^t)$ such that $N^t = \bigcup_{a \in A} \{(a, a'); a' \in N_a^t\}$. Agents receive opinions and beliefs from their neighbors only.

The distance between two opinions O and O' is measured by the absolute difference between their values:

$$d_O(O, O') = |O - O'|. \tag{2}$$

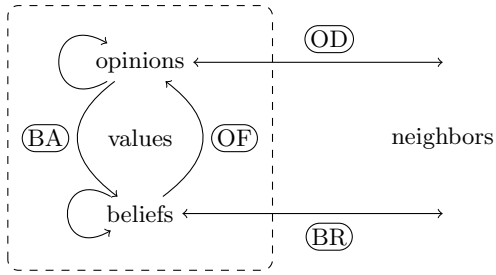

Figure 1: The four procedures (BR, OD, OF, and BA) for interactions between opinions and beliefs.

Similarly, the distance between two beliefs ϕ and ϕ' is defined as the hamming distance between their models:

$$d_B(\phi, \phi') = |\mathcal{M}(\phi) \setminus \mathcal{M}(\phi')| + |\mathcal{M}(\phi') \setminus \mathcal{M}(\phi)|. \tag{3}$$

At each time t, each agent a can determine its concordant neighbors C_a^t as those neighbors having close opinions and beliefs by:

$$C_a^t = \{a' \in N_a^t; d_O(O_a^t, O_{a'}^t) \leq \varepsilon \wedge d_B(B_a^t, B_{a'}^t) \leq \delta\}$$

such that $\varepsilon \in [0,1]$ and $\delta \in \{0, 1, \ldots, |\mathcal{M}(\top)|\}$ are threshold for opinions and beliefs, respectively.

Agents make their opinions and beliefs evolve with respect to the four procedures (OD, OF, BR, and BA) visualized in Figure 1. They can be represented as functions OD, OF, BR, and BA, respectively. These functions map $(O, B) \in \mathcal{O} \times \mathcal{L}_\mathcal{P}$ to $(O', B') \in \mathcal{O} \times \mathcal{L}_\mathcal{P}$ which means that the opinions O (resp. beliefs B) are updated to O' (resp. B') by the corresponding process. We separate these in two types of function: belief and opinion propagation (Section 5.2) and synchronization (Section 5.3).

5.2 Propagation

Agents can revise their opinions by taking other agents' opinions into account: this has been described by classical opinion dynamics models. Similarly, beliefs can be revised based on others' beliefs: this process is formalized by classical belief revision framework.

5.2.1 Opinion dynamics

In the opinion dynamics process (OD), agent a updates its current opinions using the current opinions of concordant neighbor $a' \in C_a^t$. This process can be represented as the function OD:

$$OD(O_a, B_a) = (O'_a, B_a) \tag{4}$$

such that

$$O'_a = \begin{cases} \mu O_a + (1-\mu) \times \frac{\sum_{a' \in C_a^t} O_{a'}}{|C_a^t|} & (C_a^t \neq \emptyset) \\ O_a & (C_a^t = \emptyset) \end{cases}$$

and $\mu \in [0,1]$. This procedure corresponds to the opinion updating process in [Sasahara 20] which is based on [Deffuant 00].

Table 1: Cultural values V_1 and V_2 in the case of three atoms.

p	q	r	V_1	V_2
t	t	t	7/7	0/7
t	t	f	6/7	1/7
t	f	t	5/7	2/7
t	f	f	4/7	3/7
f	t	t	3/7	4/7
f	t	f	2/7	5/7
f	f	t	1/7	6/7
f	f	f	0/7	7/7

5.2.2 Belief revision

In the belief revision process (BR), agent a updates its beliefs using the current beliefs of concordant neighbor $a' \in C_a^t$. This can be represented as the function BR with R defined as Equation 1:

$$BR(O_a, B_a) = (O_a, R(B_a, B_{x_1}, \ldots, B_{x_n})) \tag{5}$$

such that $C_a^t = \{x_1, \ldots, x_n\}$. This procedure corresponds to the belief revision operation defined in [Schwind 15].

In the definition of C_a^t, setting δ to $|\mathcal{M}(\top)|$ (resp. ε to 1) means that agents do not consider beliefs (resp. opinions) to filter their neighbors. Thus, setting $\delta = |\mathcal{M}(\top)|$ (resp. $\varepsilon = 1$) and performing OD (resp. BR) corresponds to existing model [Deffuant 00] (resp. [Schwind 15]).

5.3 Synchronization

After propagation, agents' opinions and beliefs may not correspond to each other, as shown in Section 2. This loss of coherence may be uncomfortable for agents and they may not be able to justify their opinion from their beliefs.

In order to restore coherence, we use two synchronization operations: opinion formation from beliefs (OF) and belief alignments with opinions (BA). Because opinions and beliefs are not direct translation of each others, these operations are based on the values that agents may hold.

5.3.1 Cultural values

Cultural values help agents to determine whether something is 'good' or 'bad' [Schwartz 99]. In this model, the effect of cultural values is rendered as a map V between each of the interpretations in $\mathcal{M}(\top)$ to real numbers between 0 and 1. Intuitively, the cultural values toward an interpretation I is close to 0 (resp. 1), if the agent considers I as bad (resp. good). Such values allow agents to form their opinions from beliefs and vice versa as discussed below. These values are shared among all of the agents and do not change over time.

For example, V_1 (resp. V_2) represents agents which consider the world which is consistent with p, then q, then r, as good (resp. bad).

5.3.2 Opinion formation

In the opinion formation process, agent a revises its opinions based on its beliefs. This can be achieved by mixing current opinions with the opinions based on current beliefs. The latter is calcualted by:

$$v(B) = \frac{\sum_{m \in \mathcal{M}(R(B,\phi))} V(m)}{|\mathcal{M}(R(B,\phi))|}. \tag{6}$$

Intuitively, this means that an agent forms its opinions based on its beliefs B by imagining that if the topic ϕ were true under B and evaluating how better the state, i.e., $R(B, \phi)$, is or not. Note that the divisor of $v(B)$ is never 0 since R returns consistent beliefs [Schwind 15].

Based on the function v defined above, this process can be represented as the function $OF(O, B)$:

$$OF(O_a, B_a) = (\alpha O_a + (1 - \alpha) \times v(B_a), B_a) \tag{7}$$

such that $\alpha \in [0, 1]$.

5.3.3 Belief alignment

In the belief alignment process, agents revise their beliefs based on their opinions. This can be achieved by filtering beliefs in two steps. First, agents select the beliefs $B \in \mathcal{L}_{\mathcal{P}}$ which minimize the distance between their current opinions O and the opinions based on beliefs from the function v (Equation 6). Second, agents choose beliefs which minimize the distance between their current beliefs B and the beliefs selected in the previous step.

Formally, this process can be defined as the function BA:

$$BA(O_a, B_a) = (O_a, \operatorname*{argmin}_{\phi \in \mathcal{L}_{\mathcal{P}}(O_a)} d_B(\phi, B_a)) \tag{8}$$

such that

$$\mathcal{L}_{\mathcal{P}}(O) = \operatorname*{argmin}_{\phi \in \mathcal{L}_{\mathcal{P}}} d_O(O, v(\phi)).$$

If there are several beliefs which minimize $d_B(\phi, B_a)$ in Equation 8, the aligned result is defined as one of them chosen randomly.

5.4 Protocol

Agents' behavior follows a procedure using the four operations adapted from [Sasahara 20]. At time t, the set $A^t \subseteq A$ of activated agents is determined by drawing agents in A with probability p_{activate}. Each non active agent ($a \notin A^t$) does nothing at time t; active agents will update their opinions and beliefs, and possibly their set of neighbors.

First, agent $a \in A^t$ receive its neighbors' opinions and beliefs with which it updates its own opinions O_a^t and beliefs B_a^t to get new opinions O_a^{t+1} and beliefs B_a^{t+1}. We distinguish three experimental processes to test our hypotheses: (odonly) (resp. (bronly)) only performs opinion dynamics (resp. belief revision); (both) performs both operations and apply opinion formation and belief alignment to synchronize them:

$$(O_a^{t+1}, B_a^{t+1}) = OD(O_a^t, B_a^t), \tag{odonly}$$

$$(O_a^{t+1}, B_a^{t+1}) = BR(O_a^t, B_a^t), \tag{bronly}$$

$$(O_a^{t+1}, B_a^{t+1}) = BA(OF(BR(OD(O_a^t, B_a^t)))). \tag{both}$$

All functions are executed synchronously among the set of agents.

Second, with the probability p_{rewrite}, agent $a \in A^t$ tries to modify its set of neighbors. Firstly, a chooses randomly one agent among its non concordant neighbors $a' \in N_a^t \setminus C_a^t$. Secondly, a chooses randomly one agent $a'' \in A \setminus (N_a^t \cup \{a\})$ among the non connected agents. If a can choose both of a' and a'', a updates the set of its neighbors as follows:

$$N_a^{t+1} = (N_a^t \setminus \{a'\}) \cup \{a''\}.$$

Otherwise, the set of neighbors is not updated, i.e., $N_a^{t+1} = N_a^t$. Note that for each $a \in A$, $|N_a^t|$ is constant over time t.

6. Experimental settings

Based on the model defined in Section 5, experiments to test the hypotheses of Section 4 are performed. In this section, parameters and measures for them are defined.

6.1 Parameters

The fixed parameters are as follows:

- $|A| = 100$ and $|N^0| = 400$;
- $\mathcal{P} = \{p, q, r\}$;
- topic: $\phi \equiv p$;
- $p_{\text{activate}} = p_{\text{rewrite}} = 0.5$;
- $\mu = \alpha = 0.5$;
- number of iterations: 500.

There parameters are defined so that experiments are performed under the same conditions as in [Sasahara 20].

The controled parameters are as follows:

- how agents revise their opinions and beliefs (either Equation (odonly), (bronly), or (both));
- $\varepsilon \in \{0.05, 0.1, \ldots, 0.5\} \cup \{1.0\}$;
- $\delta \in \{1, 2, \ldots, 8\}$;
- cultural values: V_1 and V_2 defined in Table 1;
- randomly generated initial states (opinions, beliefs, and networks): 3 patterns;
- seeds to manage randomness: 5 numbers.

ε is defined so that these values correspond to [Sasahara 20, Figure 4]. In [Sasahara 20, Figure 4], ε is equal to either $0.1, 0.2, \ldots, 1.0$. Here, the values for the threshold for opinions are exactly the half of what is used in [Sasahara 20] which uses the opinion space $[-1, 1]$ while for this model it is $[0, 1]$. The maximal value for δ is $|\mathcal{M}(\top)|$ with three atomic propositions.

To reduce the total number of experiments, we have performed two sets of experiments to verify **H2**: the first experiment uses fixed δ and ε varies, and the second one is the contrary. When δ (resp. ε) is fixed for **H2**, the value is taken from the experiment for **H1** (resp. **H0**) such that the average number of communities is 2, following [Sasahara 20]. Thus, the first set contains

$$\underbrace{10}_{\varepsilon} \times \underbrace{5}_{\text{seed}} \times \underbrace{3}_{\text{initial states}} \times \underbrace{2}_{\text{values}} = 300 \text{ experiments}$$

while the latter includes

$$\underbrace{7}_{\delta} \times \underbrace{5}_{\text{seed}} \times \underbrace{3}_{\text{initial states}} \times \underbrace{2}_{\text{values}} = 210 \text{ experiments.}$$

6.2 Initialization

The initial network G^0 (more specifically the initial set N^0 of edges), the initial opinions $O_1^0, \ldots, O_{|A|}^0$ and the initial beliefs $B_1^0, \ldots, B_{|A|}^0$ are generated randomly so that G^0 is connected and $B_a^0 \not\equiv \bot$ for all $a \in A$. To achieve this, N^0 is generated as follows: firstly for all $a \in A$, a' is taken from $A \setminus \{a\}$ randomly, which yields the edge (a, a'). Other $|N^0| - |A|$ edges $(x, y) \in A \times A$ are defined by generating x and y randomly so that there are no multiple edges and $x \neq y$ holds. Before starting the propagation, beliefs and opinions are synchronized first by performing (BA) (see Section 5.3). So agents start in a coherent state. Here, (OF) is not performed so that the experiments properly generalize those of [Sasahara 20].

6.3 Measures

While processing the experiments, the following measures are recorded.

6.3.1 Maximal distance measures mdb and mdo

To assess how *close* the opinions and beliefs of agents are, we use a maximal distance measure. The measure **mdb** (resp. **mdo**) refers the maximal distance between two agents' beliefs (resp. opinions) using distances defined in Equation 3 (resp. Equation 2).

6.3.2 Number of community measures ncb and nco

In order to identify how much *polarization* there is in the network, it is necessary to know the number of different communities sharing the same beliefs or opinions. The set A of agents can be partitioned in a set of communities based on their opinions or beliefs.

The belief measure **ncb** is based on the equivalence relation $\approx_{\text{B}}^t$ over A defined as:

$$a \approx_{\text{B}}^t a' \text{ iff } B_a^t \equiv B_{a'}^t,$$

which induces a partition of the set of agents $A/_{\approx_{\text{B}}^t}$. Thus, this measure is defined as $\textbf{ncb}^t = |A/_{\approx_{\text{B}}^t}|$.

The opinion measure **nco** is defined by hierarchical clustering using the distance between their opinions. The clusters are defined as the maximal subset of agents based on the clustering algorithm with the complete linkage where the differences of opinions of the members are equal to or less than 0.1. The measure **nco** is the size of this partition (number of clusters).

6.3.3 Diverse neighborhood measures anb and ano

The extreme consequence of echo chambers is network *segregation*, people connected only to people of their community. This can be measured by how much the neighbors of a community are diverse. The measure **anb** (resp. **ano**) is the average of discordant neighbors in each community based on beliefs (resp. opinions). They are defined as:

$$\textbf{anb}^t = \frac{\sum_{G \in A/_{\approx_{\text{B}}^t}} |\{(a, a') \in E^t; a \in G, a' \notin G\}|}{\textbf{ncb}^t}$$

$$\textbf{ano}^t = \frac{\sum_{G \in A/_{\approx_{\text{O}}^t}} |\{(a, a') \in E^t; a \in G, a' \notin G\}|}{\textbf{nco}^t}$$

If this value is 0, the network is extremely segregated since agents do not have any neighbor belonging to a different community.

6.4 Hypotheses revisitted

Based on the experimental plan and measures defined above, the hypotheses described in Section 4 can be revised as follows:

T0 With $\varepsilon \in \{0.05, 0.1, \ldots, 0.5\}$, $\delta = 8$ and Equation (odonly), the smaller ε is, the larger **mdo** and **nco** are.

T1 With $\varepsilon = 1$, $\delta \in \{1, 2, \ldots, 7\}$ and Equation (bronly), the smaller δ is, the larger **mdb** and **ncb** are.

T2 With Equation (both), **nco** (resp. **ncb**) is higher than with Equation (odonly) (resp. Equation (bronly)) and **ano** (resp. **anb**) is lower.

7. Results and discussions

In this section, we present the results and discuss them.

7.1 Hypothesis 0

The blue dashed line in Figure 2a shows the relation between the threshold for opinions (ε) and **mdo**. It shows that the larger ε is, the lower **mdo** is, thus the closer the opinions.

Similarly, the blue dashed line in Figure 2c shows the relation between ε and **nco**. It shows that the larger ε is, the lower **nco** is, thus the less the agents are polarized on their opinions.

Thus, **H0** is supported. These results show that the results reported in [Sasahara 20] can be reproduced.

7.2 Hypothesis 1

H1 is analyzed similarly to **H0**. The blue dashed lines in Figure 2b and Figure 2d show the relation between the threshold for beliefs (δ) and **mdb** and **ncb**, respectively. These figures show that the larger δ is, the smaller **mdb** and **ncb** are, thus the less polarized agents are on their beliefs which are closer.

Thus, **H1** is supported: we can observe echo chambers with agents who follow classical belief revision theories and ignore beliefs too different from theirs.

7.3 Hypothesis 2: effect of δ

From the results of the experiments with (bronly), it is observed that agents are split into two communities on average when $\varepsilon = 0.2$. Hence, according to the protocol of Section 6.1, we compare (bronly) with $\varepsilon = 1.0$ to (both) with $\varepsilon = 0.2$.

Figure 2d shows that **ncb** is higher with (both) than with (bronly), thus agents form more communities when opinions and beliefs interact. Moreover, Figure 2f shows that **anb** is lower with (both) than with (bronly), thus the network is more segregated.

Hence, echo chambers have been reinforced from the viewpoint of beliefs: opinions-beliefs interactions yields more polarization and more segregated network at the end of the iterations. **H2** is thus supported by these results.

7.4 Hypothesis 2: effect of ε

From the results of the experiments with (odonly), it is observed that agents are split into two communities on average when $\delta = 4$. Hence, according to the protocol of Section 6.1, we compare (odonly) with $\delta = 1$ to (both) with $\delta = 4$.

The red solid line in Figure 2c shows that if $\varepsilon \geq 0.1$, then **nco** with (both) is higher than with (odonly). Thus, there are more communities when agents consider opinions and beliefs than when they consider opinions only. In 299 experiments out of 300, agents are split into more than one community after 500 iterations, i.e., the network is polarized in the end.

Moreover, the red solid line in Figure 2e shows the effect of ε on **ano** with (both). For small ε, **ano** is larger on average than what is observed with (odonly), while it is smaller on average than before for large ε. This means that the network segregation has been reduced in the former case while it has been reinforced in the latter case. Why the former case is observed is that it is possible for agents to follow other agents who are still non concordant. In this case, agents cannot improve the homogeneity of their neighbors even after they rewrite the network. This can easily happen if only a few agents are concordant with them (i.e., they belong to a small community). Such small communities can be preserved if agents in the community consider all of their neighbors in other communities as non concordant.

Hence, interacting based on beliefs as well as opinions makes agents' opinions more polarized. On the contrary, we cannot observe more segregated networks in general. This means that **H2** is not fully supported by this experiment.

8. Conclusion

Users in social networks exchange their beliefs as well as their opinions. However, existing opinion dynamics models have not considered the beliefs that agents have. In this paper, we proposed the integration of belief revision and opinion dynamics. In this framework, agents exchange and update their opinions (resp. beliefs) using opinions (resp. beliefs) from others. Moreover, they keep the coherence between opinions and beliefs using cultural values. The results from the experiments show that echo chambers have been reinforced from the viewpoint of beliefs. However, from the opinion standpoint, less echo chambers have been observed when agents are more torelant to different opinions and beliefs from theirs.

We did not consider the influence of different values and topics, which may lead to differentiated results. Investigating such effects is left for future work.

Data availability

All experiments were performed with the *SOBA* software environment[*2]. Settings, output and data analysis notebooks are made available at [20241209-BROD], [20241210-BROD], [20241211-BROD] and [20241212-BROD].

*2 `https://github.com/tsukuba-mas/soba`

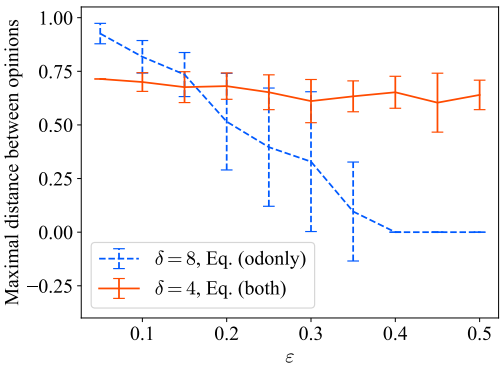

(a) The relation between ε and **mdo**.

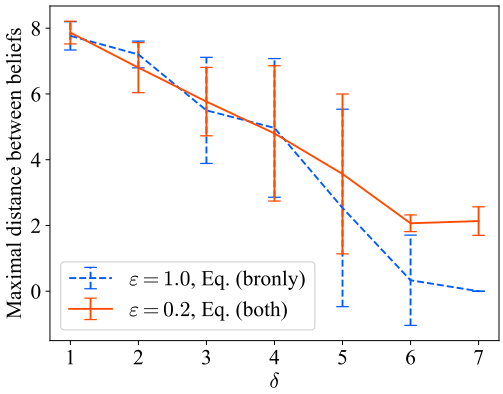

(b) The relation between δ and **mdb**.

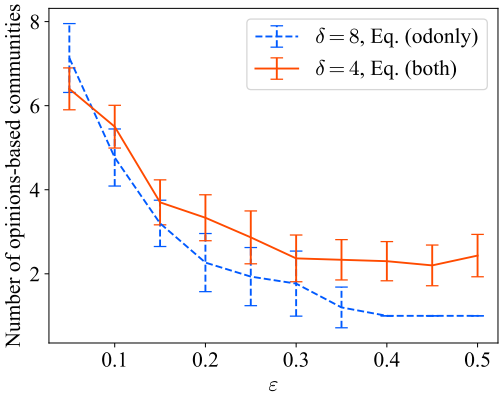

(c) The relation between ε and **nco**.

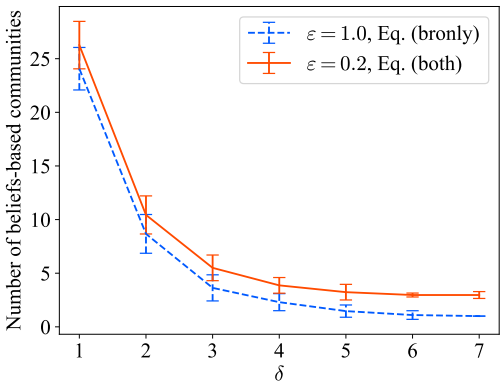

(d) The relation between δ and **ncb**.

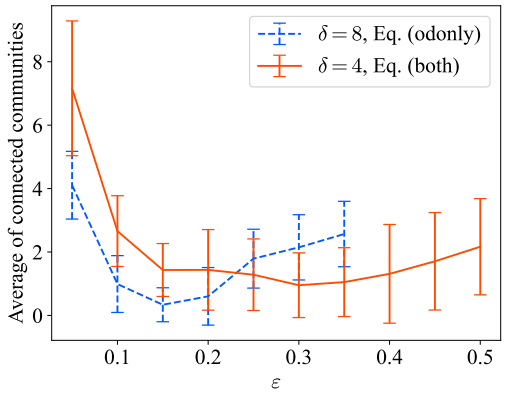

(e) The relation between ε and **ano**.

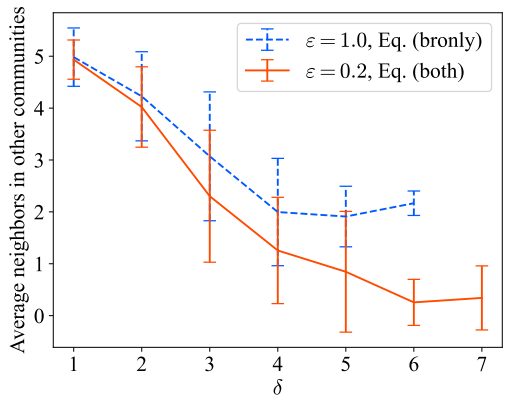

(f) The relation between δ and **anb**.

Figure 2: The difference of behavior between beliefs (resp. opinions) only (dashed blue) and when they interact (plain red). In the bottom plots, experiments which terminate with a single community are excluded.

Acknowledgments

JE has been partially supported by the MIAI Knowledge communication and evolution chair (ANR-19-P3IA-0003) and the Kyoukan PAI Project (20 006913 01 9183) supported by the Région Auvergne Rhône-Alpes. HK has been partially supported by the École Française de l'Intelligence Artificielle - MIAI Institute (Efelia, ANR-22-CMAS-0005).

Ethical Statement

The authors declare no competing interests.

The authors are not aware of any ethical issues with this work.

References

[Botan 19] Botan, S., Grandi, U., and Perrussel, L.: Multi-Issue Opinion Diffusion under Constraints, in *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 828–836, Montréal (CA) (2019)

[Cinus 22] Cinus, F., Minici, M., Monti, C., and Bonchi, F.: The Effect of People Recommenders on Echo Chambers and Polarization, *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 16, p. 90–101 (2022)

[Cossard 20] Cossard, A., De Francisci Morales, G., Kalimeri, K., Mejova, Y., Paolotti, D., and Starnini, M.: Falling into the Echo Chamber: The Italian Vaccination Debate on Twitter, in *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 14, p. 130–140 (2020)

[Deffuant 00] Deffuant, G., Neau, D., Amblard, F., and Weisbuch, G.: Mixing beliefs among interacting agents, *Advances in complex systems*, Vol. 3, No. 1-4, p. 87–98 (2000)

[DeGroot 74] DeGroot, M.: Reaching a consensus, *Journal of the American Statistical Association*, Vol. 69, No. 345, p. 118–121 (1974)

[Dong 18] Dong, Y., Zhan, M., Kou, G., Ding, Z., and Liang, H.: A survey on the fusion process in opinion dynamics, *Information Fusion*, Vol. 43, pp. 57–65 (2018)

[Friedkin 16] Friedkin, N., Proskurnikov, A., Tempo, R., and Parsegov, S.: Network science on belief system dynamics under logic constraints, *Science*, Vol. 354, No. 6310, pp. 321–326 (2016)

[Gallo 21] Gallo, F. R., Simari, G. I., Martinez, M. V., Santos, N. A., and Falappa, M. A.: Local Belief Dynamics in Network Knowledge Bases, *ACM Trans. Comput. Logic*, Vol. 23, No. 1 (2021)

[Gao 23] Gao, Y., Liu, F., and Gao, L.: Echo chamber effects on short video platforms, *Scientific Reports*, Vol. 13, No. 1 (2023)

[Hegselmann 02] Hegselmann, R. and Krause, U.: Opinion Dynamics and Bounded Confidence: Models, Analysis and Simulation, *Journal of artificial societies and social simulation*, Vol. 5, No. 3 (2002)

[Konieczny 02] Konieczny, S. and Pérez, R. P.: Merging Information Under Constraints: A Logical Framework, *Journal of Logic and Computation*, Vol. 12, No. 5, pp. 773–808 (2002)

[Mahmoudi 24] Mahmoudi, A., Jemielniak, D., and Ciechanowski, L.: Echo Chambers in Online Social Networks: A Systematic Literature Review, *IEEE Access*, Vol. 12, pp. 9594–9620 (2024)

[Nagura 25] Nagura, T. and Akiyama, E.: Emergence of echo chambers in social media communication involving multiple topics, *Information, Communication & Society*, pp. 1–22 (2025), Accepted for publication

[Nedić 19] Nedić, A., Olshevsky, A., and Uribe, C.: Graph-theoretic analysis of belief system dynamics under logic constraints, *Scientific reports*, Vol. 9, p. 8843 (2019)

[Ohagi 24] Ohagi, M.: Polarization of Autonomous Generative AI Agents Under Echo Chambers, in *Proceedings of the 14th Workshop on Computational Approaches to Subjectivity, Sentiment, & Social Media Analysis*, pp. 112–124 (2024)

[Sasahara 20] Sasahara, K., Chen, W., Peng, H., Ciampaglia, G. L., Flammini, A., and Menczer, F.: Social influence and unfollowing accelerate the emergence of echo chambers, *Journal of Computational Social Science*, Vol. 4, No. 1, p. 381–402 (2020)

[Schwartz 99] Schwartz, S. H.: A Theory of Cultural Values and Some Implications for Work, *Applied Psychology*, Vol. 48, No. 1, p. 23–47 (1999)

[Schwind 15] Schwind, N., Inoue, K., Bourgne, G., Konieczny, S., and Marquis, P.: Belief Revision Games, in *Proceedings of the 29th AAAI Conference, Austin (TX US)*, pp. 1590–1596 (2015)

[Schwind 24] Schwind, N., Inoue, K., Konieczny, S., and Marquis, P.: BeliefFlow: A framework for logic-based belief diffusion via iterated belief change, in *Proceedings of the 38th AAAI Conference, Vancouver (CA)*, pp. 10696–10704 (2024)

[Springsteen 24] Springsteen, J., Yeoh, W., and Christenson, D.: Algorithmic Filtering, Out-Group Stereotype, and Polarization on Social Media, in *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, p. 1782–1790 (2024)