Coherent belief and opinion propagation produces more echo chambers

Hiro Kataoka
Department of Computer Science,
University of Tsukuba
Tsukuba, Japan
kataoka@mas.cs.tsukuba.ac.jp

Jérôme Euzenat
Univ. Grenoble Alpes, Inria, CNRS,
Grenoble INP, LIG
F-38000 Grenoble, France
jerome.euzenat@inria.fr

Koji Hasebe
Department of Computer Science,
University of Tsukuba
Tsukuba, Japan
hasebe@cs.tsukuba.ac.jp

ABSTRACT

The occurrence of echo chambers is a well-known phenomenon in social networks. Echo chambers are groups of agents satisfying three properties: homogeneity, i.e., sharing the same opinion, segregation, i.e., listening primarily to agents of the same group, and reinforcement, i.e., increasing opinion similarity over time. Opinion dynamics models have been used to explain how echo chambers occur through agents updating their opinions. However, social network users also exchange beliefs supporting their opinions. Here, we aim at identifying the effect of propagating beliefs together with opinions on the formation of echo chambers. For that purpose, we complement an existing opinion dynamics model with classical belief merge operations. Beliefs and opinions are also maintained coherent guided by values that agents share. We introduce measures, applying to either beliefs and opinions, to identify echo chambers based on the three properties. This model is used to simulate opinion and belief propagation with varying initial conditions and agent tolerance for different opinions/beliefs. Results confirm the occurrence of echo chambers observed in previous opinion dynamics studies in a different setting and with a stricter measure. They also demonstrate the occurrence of echo chambers with belief propagation. Finally, they show that connecting opinions to beliefs and maintaining their coherence increases the number of echo chambers, but that they involve a reduced number of agents.

KEYWORDS

Echo chambers; Opinion dynamics; Belief revision; Value-based coherence; Multi-agent simulation

ACM Reference Format:
Hiro Kataoka, Jérôme Euzenat, and Koji Hasebe. 2026. Coherent belief and opinion propagation produces more echo chambers. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 9 pages. https://doi.org/10.65109/PQIV4263

1 INTRODUCTION

Echo chambers are well-known phenomena in social networks. They denote a situation in which a group of persons only interacts with people who have opinions close to theirs, thus reinforcing them [17]. They have unwanted effects in terms of social cohesion.

Opinion dynamics models have been applied to explain how echo chambers arise by aligning agents' opinions with those that they receive from others. For example, it is shown that, similarly to existing social networks, echo chambers in online social networks are created through ignoring diverging opinions and modifying the network to avoid them [22].

However, these models do not take into account the effect of agents' beliefs: people exchange and update their beliefs as well as their opinions through interactions. For example, beliefs can be communicated to justify opinions: "*I am for constructing tram lines* (opinion) because *I believe that this would reduce traffic jams* (justification by beliefs)." Hereafter, opinions and beliefs are different: the opinions of an agent denote how positive or negative it is toward a topic while beliefs represent what it considers as true based on logical reasoning. The way in which beliefs can propagate has been studied in the field of belief change [7].

Based on these observations, we aim at investigating the effect of introducing beliefs on echo chambers: Does it increase or decrease the number of opinion echo chambers? Do belief and opinion echo chambers coincide? To achieve this objective, we introduce a model defining interactions between opinions and beliefs. In this model, opinions and beliefs are propagated through existing models [22, 24] and a value-based mechanism is used to maintain their coherence [15, 16]. Using the proposed model, we aim at studying whether allowing agents to interact based on beliefs leads to more or less echo chambers.

Various measures have been provided to identify echo chambers in online social networks and simulations [3, 10, 12, 17]. However, such measures do not necessarily cover all of the characteristics of echo chambers. Thus, we introduce measures to capture echo chambers based on their definition in existing work. Introducing these measures allows us to provide new and more precise results about the influence of belief introduction on echo chambers.

Experiments confirm that agents form echo chambers either based on their opinions [22] or their beliefs [15]. In addition, they show that performing belief and opinion propagation without maintaining their coherence already generates more echo chambers. Finally, ensuring coherence between agents' opinions and beliefs leads to even more, albeit smaller, echo chambers.

The remainder of this paper is organized as follows. Section 2 introduces notations used in this paper. Section 3 presents work related to this problem. The new integrated model is defined in Section 4. Section 5 defines a measure to identify echo chambers and Section 6 draws hypotheses to be tested in this paper. Then, the experimental settings are described in Section 7. Section 8 presents and discusses the experimental results.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/PQIV4263

2 PRELIMINARIES

We introduce concepts relevant to both related work and ours.

Let A be a set of agents. 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 that agents form at time t is denoted as $G^t = \langle A, N^t \rangle$ such that $N^t = \bigcup_{a \in A}\{\langle a, a' \rangle; a' \in N_a^t\}$. For the sake of example, one can consider a set of agents called Alice, Bob, Celia, and David. The network they form is visualized in Figure 1 (left). Here, Bob and Celia are Alice's neighbors, i.e., $N^0_{\text{Alice}} = \{\text{Bob}, \text{Celia}\}$.

Each agent $a \in A$ also has a cognitive state $\langle O_a^t, B_a^t \rangle$ at time t made of its opinions O_a^t and beliefs B_a^t.

An agent's beliefs are a set of formulas held true by the agent. Here we consider agents' beliefs B as sets of formulas in a propositional logic language $\mathcal{L}_{\mathcal{P}}$, built from a set $\mathcal{P}$ of propositions. Such propositions may be that 'a tram be built' (noted as p), 'traffic jams are low' (q) or 'taxes are low' (r). If Alice believes that trams reduce traffic jams, her beliefs may be $B^0_{\text{Alice}} = \{p \rightarrow q\}$. Similarly, Bob may believe that 'building tram lines increases taxes ($B^0_{\text{Bob}} = \{p \rightarrow \neg r\}$)' and Celia that if taxes should be kept low, then trams cannot be built ($B^0_{\text{Celia}} = \{r \rightarrow \neg p\}$).

The belief set B can always be replaced by a single formula: the conjunction of its elements. In this work, this set of beliefs is expressed, without loss of generality, as a single formula $[B]$ in full disjunctive normal form (FDNF), i.e. each literal appearing exactly once in each disjunct. Hence, Alice's beliefs are expressed as $(\neg p \wedge \neg q \wedge r) \vee (\neg p \wedge q \wedge r) \vee (\neg p \wedge \neg q \wedge \neg r) \vee (\neg p \wedge q \wedge \neg r) \vee (p \wedge q \wedge \neg r) \vee (p \wedge q \wedge r)$.

The set of models of a belief set $B \subseteq \mathcal{L}_{\mathcal{P}}$ is denoted by $\mathcal{M}(B)$. There is a one-to-one correspondence between the disjuncts of the formula $[B]$ and the models of the belief set $\mathcal{M}(B)$. The symbol $\top$ (top) denotes the formula satisfied in by all interpretations (or the empty theory), thus $\mathcal{M}(\top)$ is the set of all possible interpretations.

Opinions are represented as numbers between 0 and 1, i.e., $O_a^t \in [0,1] \subseteq \mathbb{R}$ with respect to topics. In this paper, we consider a single topic ϕ which is a formula, i.e., $\phi \in \mathcal{L}_{\mathcal{P}}$. For instance, Alice, Bob, Celia, and David may discuss whether the city they are living in should construct tram lines or not. Hence they have opinions about the topic $\phi = p$ ('a tram be built'). Initially, Alice is in favor of the construction plan ($O^0_{\text{Alice}} = 1.0$), while Bob and Celia are less favorable ($O^0_{\text{Bob}} = 0.6$ and $O^0_{\text{Celia}} = 0.4$). Intuitively, opinions 0 and 1 represent extremely negative and positive opinions, respectively.

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

$$d_O(O, O') = |O - O'|. \tag{1}$$

Similarly, the distance between two belief sets B and B' is defined as the Hamming distance between their models:

$$d_B(B, B') = |\mathcal{M}(B) \setminus \mathcal{M}(B')| + |\mathcal{M}(B') \setminus \mathcal{M}(B)|. \tag{2}$$

Agents receive opinions and beliefs from their neighbors only. At each time t, agent a can determine its set of 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 $\delta \in \{0, 1, \dots, |\mathcal{M}(\top)|\}$ are thresholds for opinions and beliefs, respectively. They determine the tolerance of agents towards different opinions and beliefs. In the following examples, unless stated otherwise, we always consider that $\varepsilon = 0.4$ and $\delta = 3$. In this case, only Bob is concordant with Alice as $d_B(B^0_{\text{Alice}}, B^0_{\text{Bob}}) = 2 \leq \delta$ and $d_O(O^0_{\text{Alice}}, O^0_{\text{Bob}}) = 0.4 \leq \varepsilon$. On the opposite, $d_B(B^0_{\text{Alice}}, B^0_{\text{Celia}}) = 2 \leq \delta$ and $d_O(O^0_{\text{Alice}}, O^0_{\text{Celia}}) = 0.6 > \varepsilon$. Hence, $C^0_{\text{Alice}} = \{\text{Bob}\}$.

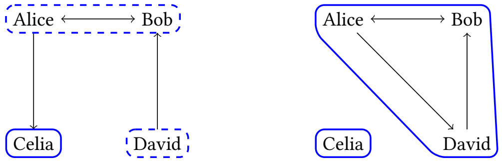

Figure 1: Visualization of a sample network before (left) and after (right) rewiring. An arrow from a person p to another person q means that p has access to q's opinions and beliefs. Blue dashed lines circumscribe strongly connected components; blue solid lines are echo chambers.

3 RELATED WORK

We first discuss the definition of echo chambers and the way to measure them. Opinion and belief propagation have been considered from the standpoint of opinion dynamics and belief change. Thus, related models in these areas are introduced in this section as well as attempts to connect them.

3.1 Echo Chambers in Social Networks

Echo chambers in social networks describe a situation in which a group of persons strengthen their viewpoints by interacting only with persons who have opinions similar to theirs. Precisely, '*An echo chamber is [...] characterized by like-minded users predominantly interacting with each other. Within these echo chambers, users express and reinforce their beliefs on specific issues, thereby amplifying their shared viewpoints*' [17]. We use the term 'echo chamber' following the definition above and not as something inherently bad [4].

Following this definition, echo chambers are groups of people satisfying three properties:

homogeneity: group members share the same or very close opinions, or beliefs;

segregation: group members mainly listen to people in the same group;

reinforcement: group members' opinions or beliefs do not become more distant over time.

There are various ways to attempt at measuring echo chambers [10]: social network analysis measures, e.g., centrality, content-based metrics comparing nodes attributes, e.g., their beliefs or opinions, homophily and polarization metrics comparing if nodes have common neighbors [12], or graph clustering and partitioning methods, also known as modularity measures [3].

Among the three properties above, homogeneity can be assessed by content-based metrics, segregation is rather complementary to homophily, but reinforcement is barely considered. In particular, [22] uses homogeneity and segregation, but not reinforcement. To identify echo chambers according to this definition, we will propose a measure taking into account the three properties.

3.2 Opinion Dynamics and Echo Chambers

How echo chambers arise has been explained using mathematical models such as opinion dynamics models. Opinion dynamics studies how agents' opinions evolve through interactions [6]. The basic idea of such models is that agents' opinions are affected by other agents' opinions. For example, if Alice has the utmost positive opinion toward the construction of tram lines, when she listens to Bob and Celia's less positive opinions toward it, she may adopt a lower opinion than before.

One classical model that can be applied to it is the *bounded confidence model* [5, 11]. In [5], two agents a and a' exchange their opinions asynchronously. They update their opinions if their difference, i.e., $d_O(O_a^t, O_{a'}^t)$, is equal to or less than a threshold $\varepsilon \in \mathbb{R}$. The updating rule can be described as:

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

such that μ is a constant between 0 and 0.5 [6]. For instance, if Alice and Bob exchange their opinions, with $\mu = 0.5$, $\varepsilon = 0.4$ and the initial opinions previously mentioned, $O^1_{\text{Alice}} = 1 + 0.5 \times (0.6 - 1) = 0.8$ and $O^1_{\text{Bob}} = 0.6 + 0.5 \times (1 - 0.6) = 0.8$. Based on this model, it is shown that low tolerance for dissimilar opinions, i.e., a small ε, prevents agents' opinions from converging.

Nowadays, opinion dynamics models based on bounded confidence have been proposed to explain how echo chambers arise [2, 22, 26]. A model with opinion threshold, timelines, and network rewiring has been proposed to model online social networks, and particularly Facebook, realistically [22]. In this model, each agent follows at least another agent. At each iteration, one agent (a) is chosen randomly. Then, agent a updates its opinions and then can modify N_a^t and submit new posts or reposts probabilistically. The updating rule is a variation of the bounded confidence model involving all the agent's concordant neighbors:

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

Assume that $\varepsilon = 0.4$, $\delta = |\mathcal{M}(\top)|$ (so that beliefs do not matter) and $\mu = 0.5$ and only Bob is concordant with Alice, then as before, O^1_{Alice} is $0.5 \cdot 1.0 + (1 - 0.5) \cdot 0.6 = 0.8$.

It is observed that the smaller the ε threshold, the more echo chambers are created. Moreover, it is reported that network rewiring, i.e., unfollowing and following, accelerates their appearance.

With the advance of large language models (LLMs), some models rely on LLM-based agents which explain reasons why they have such opinions [20, 28]. However, how agents' beliefs are changed over iterations has not been analyzed.

3.3 Distributed Belief Change

In the field of belief change [7], many operators describe how to revise or merge beliefs. For example, assume that Alice believes that 'if tram lines are built, then there will be less traffic jams ($p \longrightarrow q$)' and Bob, believes that 'if tram lines are built, then taxes will increase ($p \longrightarrow \neg r$)'. Then, after Alice listens to Bob, she may believe that the construction leads to less traffic jams *and higher taxes* by taking Bob's beliefs into account. However, little work considers belief propagation in a network.

Belief revision games (BRG) [24] define procedures, similar to those of opinion dynamics, for merging beliefs. Such an operation, R, is defined as[1]:

$$R(\mathbf{B}) = \left[\operatorname*{argmin}_{m \in \mathcal{M}(\top)} \sum_{B \in \mathbf{B}} \min_{m' \in \mathcal{M}(B)} d^H(m, m') \right] \tag{3}$$

such that $\mathbf{B}$ is a set of belief sets, those from the neighbors, to be merged and $\mathcal{M}(\top)$ is the set of all the interpretations on $\mathcal{P}$. R returns consistent beliefs [24]. For example, Alice can alter her beliefs by revising her current beliefs by Bob's beliefs. Hence, her new beliefs are $R(B^0_{\text{Alice}}, B^0_{\text{Bob}}) = \{p \longrightarrow (q \wedge \neg r)\}$ or $(p \wedge q \wedge \neg r) \vee (\neg p \wedge q \wedge r) \vee (\neg p \wedge q \wedge \neg r) \vee (\neg p \wedge \neg q \wedge r) \vee (\neg p \wedge \neg q \wedge \neg r)$ in FDNF.

BeliefFlow [25] is another framework which enables beliefs to be updated asynchronously. In this framework, agents reach a consensus if the network they form is strongly connected. None of these procedures allow agents to modify the network. Network rewiring has been considered in the network knowledge base approach on a simpler belief description language [9]. These work have not considered the impact of belief change on the formation of echo chambers, which are investigated here.

3.4 Models Evolving Opinions and Beliefs

In spite of socio-psychological evidence of their relation [21, 29], there has been few models explicitly connecting changing opinions and beliefs as defined in opinion dynamics and belief propagation. Some models constrain Boolean public opinions with common integrity constraints expressed logically [1]. Other work has used an influence matrix between opinions in order to enforce their dependencies [8, 19]. However, these constraints do not have the status of beliefs and do not change.

In [27], the distinction is made between beliefs and experience. Experience is modeled as data incompletely sampling the environment. From samples, agents learn beliefs, under the form of a classifier based on feature correlation. Then they explicitly exchange both data and classifier with their neighbors. They adopt neighbor classifiers only if it is more accurate on their own experience. This procedure is called 'epistemic rationality', meaning that agents try to maintain as much coherence as possible between their beliefs and their experience. Although this work does not deal with opinions, only belief and data, it provides the notion of internal coherence of agents' beliefs and experience.

We have introduced a way to combine and synchronise opinions and beliefs [16]. This has been extended in [15] by using confidence-bounded propagation and network rewiring in order to reproduce results of [22]. It shows that the connection between the two processes 'reinforces' echo chambers based on global measures such as maximal distance between belief/opinion, number of 'communities' (disregarding the network structure) and average number of discordant neighbors [15]. Although the second may be considered a measure of segregation, as in [22], reinforcement is completely left aside. Moreover, this work does not introduce a precise definition of echo chambers.

---

[1] d^H is the Hamming distance between two interpretations, i.e., $d^H(m, m') = |\{p \in \mathcal{P}; (m \models p \wedge m' \not\models p) \vee (m \not\models p \wedge m' \models p)\}|$.

Here, we refine this model and introduce a precise definition of echo chambers so that it is possible to count them and measure properties respective to these echo chambers instead of global properties. This enables new observations on a comparable experiment and more precision in the results. Moreover, we are able to contrast the results obtained when agents synchronize their beliefs and opinions with respect to when they do not do it.

4 A MODEL FOR OPINION AND BELIEF PROPAGATION AND INTEGRATION

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

Agents change their opinions and beliefs with respect to four procedures (*OD*, *OF*, *BR*, and *BA*). These functions map an agent's cognitive state to a new state. We separate them in two types of functions: belief and opinion propagation (Section 4.1) and coherence restoration (Section 4.2).

4.1 Propagation

Agents can revise their opinions by taking other agents' opinions into account: this corresponds to classical opinion dynamics models. Similarly, beliefs can be revised based on others' beliefs: this process is formalized by classical belief revision. Hence, we simply use existing procedures.

4.1.1 Opinion Dynamics. In the opinion dynamics process (*OD*), agents update their current opinions using the current opinions of their concordant neighbors. Thus, *OD* is defined as:

$$OD(\langle O_a, B_a\rangle) = \begin{cases} \left\langle \mu O_a + (1-\mu)\frac{\sum_{a' \in C_a^t} O_{a'}}{|C_a^t|}, B_a \right\rangle & \text{if } C_a^t \neq \emptyset \\ \langle O_a, B_a \rangle & \text{if } C_a^t = \emptyset \end{cases} \quad (4)$$

such that $\mu \in [0,1]$. This procedure corresponds to the opinion updating process in [22] which is based on the bounded confidence model (Section 3.2). Hence, if Alice performs *OD* as first operation, her cognitive state becomes $\langle 0.8, B^0_{\text{Alice}}\rangle$.

4.1.2 Belief Revision. In the belief revision process (*BR*), agents update their beliefs using the current beliefs of their concordant neighbors. Thus, *BR* is defined as:

$$BR(\langle O_a, B_a\rangle) = \langle O_a, R(B_a, B_{a_1}, \ldots, B_{a_n})\rangle \quad (5)$$

such that $C_a^t = \{a_1, \ldots, a_n\}$ are a's concordant neighbors. R is the operation defined in Section 3.3, chosen as the most regular operation among those of [24]. Accordingly, if Alice updates her beliefs with concordant agents' beliefs (here Bob's beliefs) following *BR*, her state would become $\langle O^0_{\text{Alice}}, \{p \to (q \wedge \neg r)\}\rangle$.

Setting $\delta = |\mathcal{M}(\top)|$ (resp. $\varepsilon = 1$), in the definition of C_a^t, means that agents do not ignore their neighbors based on beliefs (resp. opinions), thus performing *OD* (resp. *BR*) corresponds to the bounded confidence model [5] (resp. BRG [24]).

4.2 Coherence Between Opinions and Beliefs

After propagation, agents' opinions and beliefs may not correspond to each other. For example, after Alice listens to Bob's opinion, she may have a less positive opinion toward the construction of tram lines, but this did not affect her beliefs. Conversely, she may have changed her beliefs from Bob's without changing her opinions. Hence depending on whether opinions (*OD*) or beliefs (*BR*) have been propagated, she may found herself in a state, $\langle 0.8, B^0_{\text{Alice}}\rangle$ or $\langle O^0_{\text{Alice}}, \{p \to q \wedge \neg r\}\rangle$, which may need to be adjusted. This loss of coherence may be uncomfortable for agents as they may not be able to justify their opinions from their beliefs.

Table 1: A score V is provided to each Boolean assignment of three atoms.

p	t	t	t	t	f	f	f	f
q	t	t	f	f	t	t	f	f
r	t	f	t	f	t	f	t	f
V	$\frac{7}{7}$	$\frac{6}{7}$	$\frac{5}{7}$	$\frac{4}{7}$	$\frac{3}{7}$	$\frac{2}{7}$	$\frac{1}{7}$	$\frac{0}{7}$

In order to restore coherence, we use two synchronization operations: opinion formation from beliefs (*OF*) and belief alignment with opinions (*BA*). The former can be used to adapt opinions when beliefs have changed and the latter to adapt beliefs to new opinions. Because opinions and beliefs are not direct translations of each other, these operations are based on *values* that agents may hold.

4.2.1 Cultural Values. Cultural values help agents to determine whether something is 'good' or 'bad' [18, 23]. They are used here to determine preferences between different possible states of the world. This is rendered as a map V from 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). For example, Table 1 represents values V which consider worlds consistent with p, then q, then r, as good. In the example considered so far, these values mean that agents prefer to build tram lines more than reducing traffic jams and reducing traffic jams more than reducing taxes.

Such values allow agents to form their opinions from beliefs and vice versa as discussed below.

4.2.2 Opinion Formation. In the opinion formation process, an agent revises its opinions based on its beliefs. For example, Alice may want to lower her opinions towards trams as she has altered her beliefs after Bob's. This consists first in generating an opinion from the current beliefs (ϕ is the topic on which to form an opinion):

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

This means that an agent forms its opinion about the topic ϕ from its beliefs B based on how valuable the models of the revision of B by ϕ are. The divisor of $v(B)$ is never 0 since R returns consistent beliefs [24]. For example, after Alice has processed *BR* with Bob, her cognitive state is $\langle 1.0, \{p \to (q \wedge \neg r)\}\rangle$ (see Section 4.1.2). To adapt her opinions towards $\phi = p$ from her new beliefs, she first revises her current beliefs by ϕ: $R(p \to (q \wedge \neg r), \phi) = p \wedge q \wedge \neg r$. Assume that her values V are those defined in Table 1, she can compute the mean of the values $V(m)$ where $m \in \mathcal{M}(R(p \to (q \wedge \neg r), \phi))$, hence $v(p \to (q \wedge \neg r)) = \frac{6}{7}$.

The generated opinion is then aggregated to the agent's current opinion. The function OF is thus defined as:

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

with $\alpha \in [0,1]$ controlling the inertia to the changing process. This is a simple way to balance inertia and willingness to change opinions. For example, if Alice had opinion 1.0 toward ϕ and $\alpha = \frac{1}{2}$, i.e., her new opinion is affected by her current opinion and beliefs equally. Then, she could obtain the new opinion as: $0.5 \cdot 1.0 + (1 - 0.5) \cdot \frac{6}{7} = \frac{13}{14} \approx 0.93$. Hence, her new cognitive state is $\langle \frac{13}{14}, \{p \rightarrow (q \wedge \neg r)\}\rangle$. This result means that she considers the topic less positively than before.

4.2.3 Belief Alignment. In the belief alignment process, agents revise their beliefs based on their opinions. For example, after Alice has changed her opinion as influenced by her neighbors, she may want to alter her beliefs to align with her opinions. This is achieved by filtering beliefs in two steps. First, agents select the beliefs which minimize the distance between their current opinions and the opinions based on beliefs from the function v of Equation 6. Second, agents adopt those consistent beliefs which minimize the distance between their current beliefs and those beliefs selected in the previous step. This assumes that agents only adopt consistent beliefs.

Thus, the function BA is defined as:

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

$$\text{such that} \qquad \mathcal{L}_{\mathcal{P}}(O) = \operatorname*{argmin}_{B \in \mathrm{FDNF}^*(\mathcal{L}_{\mathcal{P}})} d_O(O, v(B))$$

and $\mathrm{FDNF}^*(\mathcal{L}_{\mathcal{P}})$ is the set of consistent FDNF formulas (excluding the empty disjunction). If several beliefs minimize $d_B(B, B_a)$ in Equation 8, the aligned result is chosen randomly among them.

For instance, after having processed OD (Section 4.1.1), Alice's state is $\langle 0.8, \{p \rightarrow q\}\rangle$. To align her beliefs to her new opinions, Alice first considers which beliefs are the best to support her opinions, i.e., which beliefs minimize the distance between her current opinion and the opinion which would be formed from her beliefs ($\mathcal{L}_{\mathcal{P}}(0.8)$). Then, she picks one of the closest beliefs $B \in \mathcal{L}_{\mathcal{P}}(0.8)$ to her current beliefs B^0_{Alice} as possible. In this case, $B = \{p \wedge r \rightarrow q\}$ (in FDNF: $(p \wedge q \wedge r) \vee (p \wedge q \wedge \neg r) \vee (p \wedge \neg q \wedge \neg r) \vee (\neg p \wedge q \wedge r) \vee (\neg p \wedge q \wedge \neg r) \vee (\neg p \wedge \neg q \wedge r) \vee (\neg p \wedge \neg q \wedge \neg r)$) minimizes the distance to B^0_{Alice}, as $d_B(B, B^0_{\text{Alice}}) = 1$. Hence, when Alice performs belief alignment, her new cognitive state is $\langle 0.8, B\rangle$.

4.3 Protocol

Agents behave following a procedure, adapted from [22], combining these four operations. 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 update their opinions and beliefs, and possibly their neighbors.

We distinguish four experimental processes to test our hypotheses: (odonly) (resp. (bronly)) only performs opinion dynamics (resp. belief revision); (both) performs both operations and (coherence) applies, in addition, opinion formation and belief alignment to synchronize them:

$$\langle O^{t+1}_a, B^{t+1}_a\rangle = OD(\ \langle O^t_a, B^t_a\rangle\ ), \tag{odonly}$$

$$\langle O^{t+1}_a, B^{t+1}_a\rangle = BR(\ \langle O^t_a, B^t_a\rangle\ ), \tag{bronly}$$

$$\langle O^{t+1}_a, B^{t+1}_a\rangle = BR(\ OD(\ \langle O^t_a, B^t_a\rangle\ )\ ), \tag{both}$$

$$\langle O^{t+1}_a, B^{t+1}_a\rangle = OF(\ BR(\ BA(\ OD(\ \langle O^t_a, B^t_a\rangle\ )\ )\ )\ ). \tag{coherence}$$

All functions are executed synchronously among active agents.

Then, as in [22], with the probability p_{rewire}, agent $a \in A^t$ tries to modify its set of neighbors. Intuitively, this is achieved by removing from N^t_a one agent a' that is not concordant and adding one agent a'' that is not a neighbor or a itself to it. More formally, a considers the sets $F^t_a \subseteq A$ and $U^t_a \subseteq A$ of agents that can be followed and unfollowed at time t, respectively. They are defined as:

$$U^t_a = N^t_a \setminus C^t_a \text{ and } F^t_a = \{a' \in A; a \neq a' \wedge a' \notin N^t_a\}.$$

If the two sets are non empty, a can choose the unfollowed agent $a' \in U^t_a$ and the followed agent $a'' \in F^t_a$ randomly, and update its set of neighbors as follows:

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

Otherwise, the set of neighbors is not updated, i.e., $N^{t+1}_a = N^t_a$. This entails that for each $a \in A$, the size of N^t_a remains constant over t. For example, when Alice performs the rewiring, the unfollowed agent a' is Celia because $U^0_{\text{Alice}} = \{\text{Celia}\}$. Moreover, in the network presented in Figure 1 (left), David is the only agent who is not a neighbor of Alice, i.e., $F^0_{\text{Alice}} = \{\text{David}\}$. Hence, Alice can update her neighbors as:

$$N^1_{\text{Alice}} = (N^0_{\text{Alice}} \setminus \{\text{Celia}\}) \cup \{\text{David}\} = \{\text{Bob}, \text{David}\}.$$

Figure 1 (right) shows the network after the rewiring.

5 MEASURES

In [22], the emergence of echo chambers is shown visually using plots of the evolution of message diversity, agents' opinions, and the network of agents over time. However, these plots are generated from results of one single experiment.

In the same paper, the effects of ε on the number of final opinion peaks and final opinion distance are presented based on 20 experiments which share the same parameters. These measures show how opinions are polarized and the network segregated. Although polarization may be a measure of homogeneity, reinforcement is ignored, so this should not allow us to conclude on echo chambers. Results have been confirmed in [15] but still with measures not explicitly identifying echo chambers.

The way to identify echo chambers is to partition the network into groups and to count those satisfying the properties of homogeneity, segregation and reinforcement (see Section 3.1). Of course, it is not practicable to test all partitions, thus we propose to use strongly connected components to identify echo chambers. Basing measurements on strongly connected components make sense because: (1) echo chambers are expected to be connected, (2) they are maximal, and (3) reinforcement is usually obtained through feedback, which needs strong connection. Nevertheless, strongly connected components may be replaced by other structures, partitions or not. The effect of this choice is left for further investigation.

Table 2: Configuration of the parameters. $\mathbf{eo}_X^t$ and $\mathbf{eb}_X^t$ mean the measure $\mathbf{eo}^t$ and $\mathbf{eb}^t$ obtained in the setting X, respectively.

Hypothesis	H0	H1	H2	H3
Setting	S0	S1	S2	S3
protocol	(odonly)	(bronly)	(both)	(coherence)
ε	$\{0.05, 0.1, \dots, 0.5\}$	1	$\{0.05, 0.1, \dots, 0.5\}$	$\{0.05, 0.1, \dots, 0.5\}$
δ	8	$\{1, 2, \dots, 7\}$	$\{1, 2, \dots, 7\}$	$\{1, 2, \dots, 7\}$
#Sims	$20 \cdot 10 = 200$	$20 \cdot 7 = 140$	$20 \cdot 10 \cdot 7 = 1400$	$20 \cdot 10 \cdot 7 = 1400$
Test	$\mathbf{eo}_{S0}^T > 0$	$\mathbf{eb}_{S1}^T > 0$	$\mathbf{eo}_{S2}^T \geq \mathbf{eo}_{S0}^T$ and $\mathbf{eb}_{S2}^T \geq \mathbf{eb}_{S1}^T$	$\mathbf{eo}_{S3}^T \geq \mathbf{eo}_{S2}^T$ and $\mathbf{eb}_{S3}^T \geq \mathbf{eb}_{S2}^T$

Hereafter, *component* denotes strongly connected components and $\mathcal{S}^t$ is the set of components at time t.

During the experiments, the following measures are recorded.

Segregation. The *segregation* of a component is measured by the ratio of the edges whose sources are in a component C and whose targets are not in C to the edges whose sources are in C:

$$L^t(C) = \frac{|\{\langle a, a' \rangle \in N^t; a \in C \land a' \notin C\}|}{|\{\langle a, a' \rangle \in N^t; a \in C\}|}.$$

Hence, the lower this measure, the more segregated the network.

Homogeneity. To assess how *homogeneous* agents in a component $C \in \mathcal{S}^t$ are, we measure the maximal distance between opinions (resp. beliefs) of agents which belong to the component:

$$M_O^t(C) = \max_{a, a' \in C} d_O(O_a^t, O_{a'}^t),$$
$$M_B^t(C) = \max_{a, a' \in C} d_B(B_a^t, B_{a'}^t).$$

The lower these measures, the more homogeneous the component.

Reinforcement. To assess if a component C satisfies *reinforcement*, we test whether its homogeneity measure (maximal distance between agents' opinions or beliefs) decreases during its lifespan. Let $[t_C, t]$ be the maximal window during which a component $C \in \mathcal{S}^t$ exists, then, the measure is defined as:

$$D_O^t(C) \equiv \forall s \in [t_C, t[, M_O^s(C) \geq M_O^{s+1}(C)$$
$$D_B^t(C) \equiv \forall s \in [t_C, t[, M_B^s(C) \geq M_B^{s+1}(C)$$

Echo chambers. Based on the above measures, we can compute the number of echo chambers at time t as those components that are segregated and whose opinions and beliefs are homogeneous and reinforcing:

$$\mathbf{eo}^t = |\{C \in \mathcal{S}^t; \overbrace{L^t(C) \leq \theta}^{\text{segregation}} \land \overbrace{M_O^t(C) \leq 10^{-5}}^{\text{homogeneity}} \land \overbrace{D_O^t(C)}^{\text{reinforcement}}\}|,$$
$$\mathbf{eb}^t = |\{C \in \mathcal{S}^t; L^t(C) \leq \theta \land M_B^t(C) = 0 \land D_B^t(C)\}|.$$

such that $\theta \in]0, 1[$. The measures are made independent of ε and δ, which vary in experiments, by choosing stricter thresholds (10^{-5} and 0) for the homogeneity property. In the following, we use $\theta = 0.5$. Here, $L^t(C) \leq 0.5$ indicates that at most half of the edges whose sources are in C, have their targets out of C. This automatically disqualifies singletons (components C such that $|C| = 1$) as echo chambers. In complement, we compute the proportion of agents involved in echo chambers.

In Figure 1 (left), there are three strongly connected components, but only Celia who does not listen to any one could be an echo chamber, as Alice, Bob and David have significantly different opinions and beliefs. After propagation and rewiring (right), Alice, Bob and David form a strongly connected component that is an echo chamber.

6 HYPOTHESES

We test the four following hypotheses taking into account the more elaborate definition of echo chambers.

First, the proposed model is defined based on that of [22]. So, if only opinions are updated, echo chambers should be observed:

H0 If opinions evolve independently from beliefs, opinion echo chambers should arise.

Similarly with respect to [15], the same should occur for beliefs:

H1 If beliefs evolve independently from opinions, belief echo chambers should arise.

Next, it is expected that independent opinion and belief propagation reinforce each other, through the double constraint put on concordant agents, and thus lead to more echo chambers:

H2 Connecting opinions and beliefs increases the number of echo chambers.

Finally, keeping opinions and beliefs coherent is expected to tighten the links between them and to induce more echo chambers than without keeping the coherence:

H3 Connecting opinions and beliefs and keeping them coherent increases the number of echo chambers.

7 EXPERIMENTAL SETTINGS

In order to test each of the four hypotheses of Section 6, we perform four experiments under specific conditions described hereafter.

We fix $\mathcal{P}$ to $\{p, q, r\}$ and the topic ϕ to p. Moreover, we fix $|A| = 100$ and $|N^0| = 400$ to follow the same conditions as those of [15, 22]. We always use $T = 5000$ iterations.

In addition, $p_{\text{active}} = p_{\text{rewire}} = 0.5$ and $\mu = 0.5$ to perform experiments under the same conditions as the experiments used to generate Figure 3 in [22]. We fix $\alpha = 0.5$ so that the effect of an agent's beliefs on its opinions is the same as the effect of other agents' opinions on it. Finally, the cultural values V presented in Table 1 is shared among all of the agents and does not change over time.

The initial state is generated by drawing randomly $|N^0|$ edges among $|A|$ agents, so that each agent has at least one neighbor.

Table 3: Overall results for the four experimental settings and the initial setting. In parentheses for components are the percentages of agents belonging to such components. Raw results are aggregated by computing the mean.

Setting	initial	S0	S1	S2	S3
stabilized experiments	–	200 (100%)	140 (100%)	1389 (99.2%)	1390 (99.3%)
components ($\lvert\mathcal{S}^T\rvert$)	3.35 (100%)	4.29 (100%)	7.48 (100%)	7.97 (100%)	8.29 (100%)
singletons	2.35 (2.35%)	2.35 (2.35%)	4.39 (4.39%)	4.25 (4.25%)	4.29 (4.29%)
opinion echo chambers ($\mathbf{eo}^T$)	0.00 (0%)	1.71 (97.2%)	–	2.70 (86.8%)	3.00 (81.7%)
beliefs echo chambers ($\mathbf{eb}^T$)	0.00 (0%)	–	2.03 (79.3%)	2.52 (74.9%)	3.00 (80.4%)
Test	–	√	√	√ and √	√ and √

Agents are assigned random beliefs and opinions. We use 20 different seeds to control the randomness of both the initial network and the propagation process. Before starting interacting with each other, all agents perform $BA \circ OF$ until their opinions and beliefs become stable, i.e., each agent's beliefs do not change and the difference between opinions before and after the operation is less than or equal to 10^{-5}. Thus, agents are in a coherent state at the beginning.

δ is taken from the set $\{1, 2, \ldots, |\mathcal{M}(\top)|\} = \{1, 2, \ldots, 8\}$. ε is taken from the set $\{0.05, 0.1, \ldots, 0.5\}$ so that the experiments correspond to those performed in [22] and [15].

Table 2 relates controlled parameters to hypotheses. The different experiment sizes (#Sims) are the consequence of the controlled variables with respect to the hypotheses: when a parameter is frozen, all values lead to the same result.

8 RESULTS AND DISCUSSION

Results of the four experiments are summarized in Table 3. Thereby, 'stabilized experiments' is the number of experiments in which agents' opinions and beliefs did not change during the last 500 iterations, i.e., $B_a^t \equiv B_a^T$ and $|O_a^t - O_a^T| \leq 10^{-5}$. Although agents' cognitive states did not change, the network itself may change due to network rewiring, albeit without affecting agent states.

Hypothesis H0. In each simulation with S0, $\mathbf{eo}^T > 0$, i.e., at least one echo chamber is created. Thus, the results support Hypothesis H0. In addition, all components (not only echo chambers) satisfy reinforcement. Figure 2 (left plot, solid green curve) shows the decrease of echo chambers with the increase of tolerance (higher ε). Hence, although we applied a stronger measure, results confirm those of [22].

Figure 3 (a) shows what happens on a specific case: a single large component gathers most agents with the same opinion. A few singleton components listen to the echo chamber. The echo chamber effect is huge in this case, as in many, as all agents have the same opinion.

Hypothesis H1. In each simulation with S1, but 8 out of 140 (94%), $\mathbf{eb}^T > 0$. These results show that we can observe belief echo chambers in most of the cases. Thus, Hypothesis H1 is supported.

As for opinions, Figure 2 (right plot, solid orange curve) shows the number of echo chambers decrease with the increase of tolerance (higher δ).

Figure 3 (b) shows the network scattered into many smaller components, with more echo chambers and more diverse beliefs.

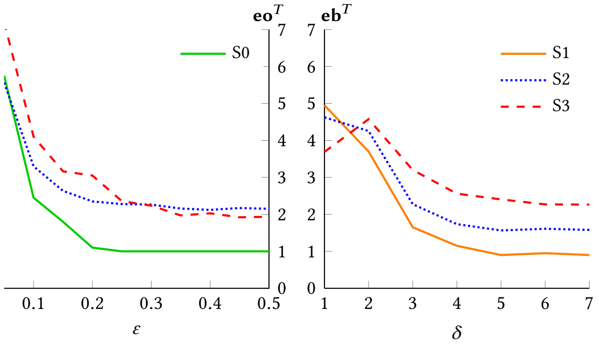

Figure 2: Effect of ε, resp. δ, on the number of opinion, resp. belief, echo chambers. More tolerant agents lead to less echo chambers, but connecting opinions and beliefs leads to more echo chambers.

Hypothesis H2. Table 3 shows that, on average, connecting opinions and beliefs (S2) increases the number of opinion (S0) and belief (S1) echo chambers. This supports Hypothesis H2. In counterpart, the number of agents involved in these echo chambers, especially with respect to opinions, has decreased. Figure 3 (c) shows that this is not only due to singleton components.

An interesting remark is that, the number of echo chambers increases without directly synchronizing opinions and beliefs. This is explained by the rewiring of the network, which generates segregation. Segregation, because agents start with coherent beliefs and opinions, tends to perpetuate coherence. In consequence, echo chambers on both opinions and beliefs are closely related.

The echo chambers in Figure 3 (c) (and d) are indeed homogeneous in both opinions and beliefs.

Hypothesis H3.

Table 3 shows that simulations with S3 lead, on average, to more echo chambers than those with S2 involving roughly the same proportion of agents. This would support Hypothesis H3. However, opinion and belief echo chambers increase or stay stable in only 78.86% of the experiments (1104 out of 1400). There is no direct relation between S2 and S3: pairwise comparison does not indicate a systematic increase, as shown in Table 4.

Figure 3 (c) and (d) show an increase of echo chambers (6 to 7).

Table 4: Distribution of experiments with respect to the effect of maintaining coherence (S3) compared to not maintaining it (S2) on eb^T and eo^T (all other parameters being the same).

eb^T / **eo**^{T}	decreased	did not change	increased
decreased	258 (18.4%)	29 (2.07%)	3 (0.21%)
did not change	6 (0.43%)	482 (34.4%)	61 (4.36%)
increased	0	3 (0.21%)	558 (39.9%)

9 CONCLUSION

In social networks, symbolic beliefs are propagated as well as opinions. To test their influence on echo chambers, we use a model in which agents exchange and update their opinions and/or beliefs using those of their neighbors. Moreover, they maintain the coherence between their opinions and beliefs using cultural values.

We identify echo chambers based on components satisfying segregation, homogeneity and reinforcement of opinions or beliefs. Experimental results confirm those obtained with different measures: the more tolerant agents are, the less echo chambers they create. This also shows that joining opinions and beliefs generates more echo chambers. This occurs even when agents do not explicitly maintain the coherence between their opinions and beliefs.

Hence, connecting beliefs to opinions indeed affects both types of echo chambers, which consolidate each other. This does not prevent echo chambers from arising. On the contrary, there are more echo chambers, albeit smaller ones (see Figure 3), involving slightly fewer agents. Thus, it is worth asking if this strengthens or weakens echo chambers. This could be tested by evaluating how the resulting echo chambers persist in the case of perturbations, i.e. locally modifying the graph topology or agent states.

Data availability. All experiments were performed with the *SOBA* simulator [13]. The experiment notebook is openly available from [14]. It contains reference to software version, processing instructions, results, and analyses, including statistical tests.

Ethical statement. The authors declare no competing interests. The authors are not aware of any ethical issue with this work.

ACKNOWLEDGMENTS

This work was partially supported by the Fukuro INRIA-University of Tsukuba associate team and the Echo action of the NTU-UGA-UT trilateral center. HK has been partially supported by the École Française de l'Intelligence Artificielle - MIAI Institute (Efelia, ANR-22-CMAS-0005).

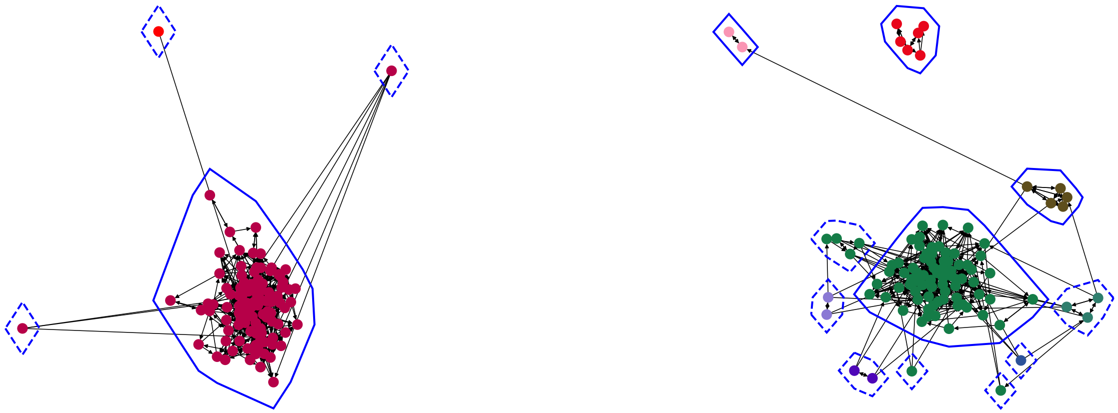

(a) Final network with S0 (Same color means same opinion).

(b) Final network with S1 (Same color means same belief).

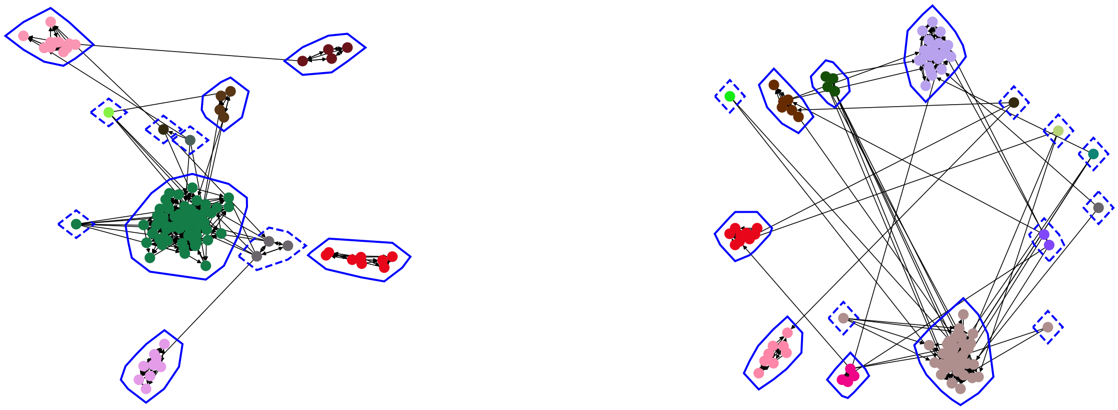

(c) Final network with S2 (Same color means same belief and opinion).

(d) Final network with S3 (Same color means same belief and opinion).

Figure 3: Networks at T from experiments using $\varepsilon = 0.2$, $\delta = 2$, and the same seed.

REFERENCES

[1] Sirin Botan, Umberto Grandi, and Laurent Perrussel. 2019. Multi-Issue Opinion Diffusion under Constraints. In *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. Montréal (CA), 828–836. https://www.ifaamas.org/Proceedings/aamas2019/pdfs/p828.pdf

[2] Federico Cinus, Marco Minici, Corrado Monti, and Francesco Bonchi. 2022. The Effect of People Recommenders on Echo Chambers and Polarization. *Proceedings of the International AAAI Conference on Web and Social Media* 16 (May 2022), 90–101. https://doi.org/10.1609/icwsm.v16i1.19275

[3] Aaron Clauset, M. E. J. Newman, and Cristopher Moore. 2004. Finding community structure in very large networks. *Phys. Rev. E* 70 (Dec 2004), 066111. Issue 6. https://doi.org/10.1103/PhysRevE.70.066111

[4] David Coady. 2024. Stop Talking about Echo Chambers and Filter Bubbles. *Educational Theory* 74, 1 (2024), 92–107. https://doi.org/10.1111/edth.12620

[5] Guillaume Deffuant, David Neau, Frédéric Amblard, and Gérard Weisbuch. 2000. Mixing beliefs among interacting agents. *Advances in complex systems* 3, 1-4 (2000), 87–98. https://doi.org/10.1142/S0219525900000078

[6] Yucheng Dong, Min Zhan, Gang Kou, Zhaogang Ding, and Haiming Liang. 2018. A survey on the fusion process in opinion dynamics. *Information Fusion* 43 (2018), 57–65. https://doi.org/10.1016/j.inffus.2017.11.009

[7] Eduardo Fermé and Sven Ove Hansson. 2018. *Belief change: introduction and overview*. Springer, Cham (CH).

[8] Noah Friedkin, Anton Proskurnikov, Roberto Tempo, and Sergey Parsegov. 2016. Network science on belief system dynamics under logic constraints. *Science* 354, 6310 (2016), 321–326. https://doi.org/10.1126/science.aag2624

[9] Fabio R. Gallo, Gerardo I. Simari, Maria Vanina Martinez, Natalia Abad Santos, and Marcelo A. Falappa. 2021. Local Belief Dynamics in Network Knowledge Bases. *ACM Trans. Comput. Logic* 23, 1, Article 4 (10 2021), 36 pages. https://doi.org/10.1145/3477394

[10] David Hartmann, Sonja Mei Wang, Lena Pohlmann, and Bettina Berendt. 2025. A systematic review of echo chamber research: comparative analysis of conceptualizations, operationalizations, and varying outcomes. *Journal of Computational Social Science* 8, 2 (April 2025). https://doi.org/10.1007/s42001-025-00381-z

[11] Rainer Hegselmann and Ulrich Krause. 2002. Opinion Dynamics and Bounded Confidence: Models, Analysis and Simulation. *Journal of artificial societies and social simulation* 5, 3 (2002). https://www.jasss.org/5/3/2.html

[12] Ruben Interian, Ruslán G. Marzo, Isela Mendoza, and Celso C. Ribeiro. 2023. Network polarization, filter bubbles, and echo chambers: an annotated review of measures and reduction methods. *International Transactions in Operational Research* 30, 6 (2023), 3122–3158. https://doi.org/10.1111/itor.13224

[13] Hiro Kataoka. 2025. SOBA simulator. https://github.com/tsukuba-mas/soba.

[14] Hiro Kataoka. 2026. Propagating and synchronizing beliefs and opinions generate more echo chambers than non synchronizing. https://sake.re/20260217-BROD/.

[15] Hiro Kataoka, Jérôme Euzenat, and Koji Hasebe. 2025. Exchanging and updating opinions and beliefs reinforces echo chambers. In *Proc. 39th annual conference of the Japanese society for artificial intelligence (JSAI), Osaka (JP).*

[16] Hiro Kataoka, Jérôme Euzenat, and Koji Hasebe. 2026. How can beliefs alter opinions? Joint opinion and belief evolution. In *Proceedings of the 18th International Conference on Agents and Artificial Intelligence (ICAART), Marbella (ES).* 1011-1022.

[17] Amin Mahmoudi, Dariusz Jemielniak, and Leon Ciechanowski. 2024. Echo Chambers in Online Social Networks: A Systematic Literature Review. *IEEE Access* 12 (2024), 9594–9620. https://doi.org/10.1109/ACCESS.2024.3353054

[18] Rijk Mercuur, Virginia Dignum, and Catholijn Jonker. 2019. The value of values and norms in social simulation. *Journal of artificial societies and social simulation* 22, 1 (2019). https://doi.org/10.18564/jasss.3929

[19] Angelina Nedić, Alex Olshevsky, and César Uribe. 2019. Graph-theoretic analysis of belief system dynamics under logic constraints. *Scientific reports* 9 (2019), 8843. https://doi.org/10.1038/s41598-019-45076-4

[20] Masaya Ohagi. 2024. Polarization of Autonomous Generative AI Agents Under Echo Chambers. In *Proceedings of the 14th Workshop on Computational Approaches to Subjectivity, Sentiment, & Social Media Analysis*. 112–124. https://aclanthology.org/2024.wassa-1.10

[21] Kerem Oktar and Tania Lombrozo. 2025. How aggregated opinions shape beliefs. *Nature Reviews Psychology* 4, 2 (Jan. 2025), 81–95. https://doi.org/10.1038/s44159-024-00398-7

[22] Kazutoshi Sasahara, Wen Chen, Hao Peng, Giovanni Luca Ciampaglia, Alessandro Flammini, and Filippo Menczer. 2020. Social influence and unfollowing accelerate the emergence of echo chambers. *Journal of Computational Social Science* 4, 1 (2020), 381–402. https://doi.org/10.1007/s42001-020-00084-7

[23] Shalom H. Schwartz. 1999. A Theory of Cultural Values and Some Implications for Work. *Applied Psychology* 48, 1 (Jan. 1999), 23–47. https://doi.org/10.1111/j.1464-0597.1999.tb00047.x

[24] Nicolas Schwind, Katsumi Inoue, Gauvain Bourgne, Sébastien Konieczny, and Pierre Marquis. 2015. Belief Revision Games. In *Proceedings of the 29th AAAI Conference, Austin (TX US)*. 1590–1596. https://doi.org/10.1609/aaai.v29i1.9415

[25] Nicolas Schwind, Katsumi Inoue, Sébastien Konieczny, and Pierre Marquis. 2024. BeliefFlow: A framework for logic-based belief diffusion via iterated belief change. In *Proceedings of the 38th AAAI Conference, Vancouver (CA)*. 10696–10704. https://doi.org/10.1609/aaai.v38i9.28941

[26] Jean Springsteen, William Yeoh, and Dino Christenson. 2024. Algorithmic Filtering, Out-Group Stereotype, and Polarization on Social Media. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems* (Auckland, New Zealand). 1782–1790.

[27] Yanni Tang, Jiamou Liu, and Wu Chen. 2021. Exchange, adopt, evolve: Modeling the spreading of opinions through cognition and interaction in a social network. *Information sciences* 551 (2021), 1–22. https://doi.org/10.1016/j.ins.2020.11.043

[28] Chenxi Wang, Zongfang Liu, Dequan Yang, and Xiuying Chen. 2024. Decoding Echo Chambers: LLM-Powered Simulations Revealing Polarization in Social Networks. arXiv:2409.19338 [cs.SI] https://arxiv.org/abs/2409.19338

[29] Tao Wen, Rui Zheng, Ting Wu, Zeyi Liu, Mi Zhou, Tahir Abbas Syed, Darminder Ghataoura, and Yuwang Chen. 2025. Formulating opinion dynamics from belief formation, diffusion and updating in social network group decision-making: Towards developing a holistic framework. *European Journal of Operational Research* 325, 3 (2025). https://doi.org/10.1016/j.ejor.2024.12.015