Beliefs in echo chambers*

Pranjal Bhushan† Abhinash Borah‡ Siddarth Venkatesh‡

January 7, 2025

Abstract

We introduce an interactive model of belief formation and transmission within echo chambers. Individuals in our model have subjective core beliefs, but these are not the same as the beliefs underlying their behavior. An individual's *behavioral beliefs* incorporate her core beliefs but are also influenced by the behavioral beliefs of others within her echo chamber. Therefore, echo chambers feature interactive behavioral beliefs with any individual's beliefs both being influenced by as well as influencing the beliefs of others in her echo chamber. The *echo chamber representation* of a profile of such behavioral beliefs that we propose captures the steady state of this process of interaction and influence. We show that the model is falsifiable by characterizing it based on two axioms. The first emphasizes the need for conformity when it comes to assessments about certain events, while the second highlights the possibility of differing behavioral beliefs about uncertain events and the potential for everyone to exercise influence. We also analyze when the model permits exact identification, i.e., when can an analyst draw on the profile of behavioral beliefs to uniquely identify the composition of the echo chambers in society, the core beliefs of different individuals, and the degree to which each of them is immune to influence. Further, we provide an observational learning foundation for our theory by incorporating sequential arrival of private information into the DeGroot (1974) model. We use this foundation to provide a novel theory of interactive belief updating for non-Bayesian decision-makers.

JEL codes: D01, D81, D91

Keywords: echo chambers, interactive beliefs and updating, observational learning, influence and immunity from influence, conformism, characterization and identification

*We thank Ariel Rubinstein for helpful comments. We also thank conference/seminar participants at DSE, ISI-Delhi, HEC Paris, BREW, Ahmedabad University, AMES-CSW, and the Economics Exchange at Ashoka.

†HEC Paris (email: pranjal.bhushan@hec.edu)

‡Ashoka University (emails: abhinash.borah@ashoka.edu.in, siddarthv2011@gmail.com)
Correspondence address: Department of Economics, Ashoka University. Rajiv Gandhi Education City, Sonepat. Haryana - 131029, India.

1 Introduction

Research across several disciplines has shown that *echo chambers* that form in society impact a range of important social and economic outcomes like polarization, populism, inequality, and asset prices in financial markets (Barberá 2020; Cookson et al. 2023; McCarty et al. 2006). A key reason echo chambers drive these outcomes is the beliefs that form within them and how they get transmitted. This paper looks at beliefs formed within echo chambers. Specifically, we propose a theory of interactive beliefs within echo chambers that capture how individual beliefs are formed based both on subjective assessments of uncertainty that incorporates private information as well as the influence cast by popular perceptions within one's echo chamber that encourages conformism.

Our theory builds on the following idea. We consider a society that is partitioned into a set of echo chambers. Each individual in society has some *core beliefs* over the states of the world that captures her subjective assessment of uncertainty and incorporates her private information. However, her behavior is not determined exclusively by her core beliefs. Rather, her *behavioral beliefs* also draw on the overall beliefs prevailing within her echo chamber. Specifically, we model an individual's behavioral beliefs as a weighted average of her core beliefs and the average behavioral beliefs in her echo chamber. The weight put on the former measures the degree to which she is immune from her echo chamber's influence and is an important behavioral parameter. Echo chambers, therefore, feature interactive behavioral beliefs with an individual's beliefs both being influenced by and at the same time influencing the beliefs of others in her echo chamber. The *echo chamber representation* of a profile of such behavioral beliefs that we introduce in this paper captures the steady state of this process of interaction and influence.

The substantive work that we undertake in this paper addresses several important questions about this interactive process. First, we show that our interactive model of beliefs is falsifiable by providing a precise characterization of the steady state beliefs underlying an echo chamber representation. That is, consider an analyst who has data on the behavioral beliefs of individuals in society, say, from observing their betting behavior. Our characterizing conditions establish the restrictions that this data must satisfy for her to conclude that this profile of beliefs is consistent with the echo chamber model. Second, we show the extent to which the model parameters—the echo chambers that form, each individual's core beliefs and their susceptibility to influence—can be uniquely identified from behavioral beliefs. Third, we provide an observational learning foundation for the echo chamber model. We do so by incorporating private information into a standard DeGroot (1974) learning process. Finally, we use these foundations for the echo chamber model to propose a theory of interactive belief updating that combines Bayesian updating of core beliefs with naive observational learning, where the key source of naivety stems from the way in which the process of interaction and influence results in information not being fully incorporated into updates and the failure to recognize this.

Our characterizing conditions highlight the key restrictions underlying interactive belief formation. The first condition is called *certainty conformism*. It says that for any decision maker (DM) to assign probability one to an event according to her behavioral beliefs, i.e., behave as if she is certain about it, she must have the confirmation of at least someone else who does likewise, with this process being mutual. This captures a key behavioral restriction of the echo chamber model. No DM can behave as if she is sure about an event on her own. She can do so only when there are others (her revealed friends) who do so. This therefore generates conformity when it comes to sure events. At the same time our second condition, called *subjective assessment of uncertainty*, makes the point that inter-personal interactions in an echo-chamber is not just about conformity. Rather, such interactions also incorporate the possibility of individuals having different behavioral beliefs about uncertain events from that of their friends, thus revealing the potential for mutual influence.

Together, these two conditions connect interesting themes across the literatures on social influence and decisions under risk. Keeping with the former, the axioms reiterate that social influence is not just about conformity, but also about inter-personal differences that open the scope for differential influence. When it comes to the latter, the axioms highlight the point that sure events may be viewed qualitatively differently by DMs than uncertain events. The need for complete conformity is restricted only to sure events and doesn't necessarily spill over to events about which people maintain a degree of uncertainty. This qualitative distinction between sure events and uncertain events resonates with similar distinctions that have been made in the context of theories like rank dependent utility and prospect theory, which draw on non-linear probability weighting, in especially pronounced ways around unit (and zero) probability events, and in such observed phenomenon like the certainty effect.

We provide another foundation for echo chamber beliefs by showing how they emerge as the limiting beliefs under a process of DeGroot learning, where agents sequentially receive private information. Our two characterizing conditions above, therefore, provide a fresh perspective on characterizing DeGroot learning with private information in an environment where the analyst doesn't have the opportunity to observe the entire observational learning path but potentially only the steady state beliefs that result from such a process. Of course, such a perspective makes sense only if the sequence in which private information arrives doesn't affect the steady state beliefs. This is indeed true for the observational learning process underlying the echo chamber model. This distinguishes our learning story from other adaptations of the DeGroot model to incorporate sequential arrival of private information in which steady state beliefs are not invariant to the sequence in which information arrives (e.g., Reshidi (2024)).

We use our observational learning setup to analyze how behavioral beliefs respond to a single release of private information. Beliefs respond to information through two channels. The *direct* effect arises from Bayesian updating of core beliefs. The *echo* effect is a result of naïvely accounting for changes in the behavioral beliefs of others in the echo chamber. The two effects together result in behavioral beliefs that tend to be sticky, with individuals failing to fully incorporate private information unless everyone in the echo chamber receives the same information. This stickiness results in the persistence of beliefs inspite of new evidence, increasing the susceptibility of individuals to confirmation bias. In the context of majority voting in a common value setting, this results in sub-optimal outcomes due to incorrect aggregation of information with a greater probability than in the absence of influence, while also making restrictions on the information source more stringent for the wisdom of the crowd to prevail.

The literature on echo chambers is primarily divided between examining the process of segregation into homogeneous groups and the transmission of beliefs and biases through these groups thereafter. Segregation into various groups is attributed to a variety of reasons, including economic, social, and cultural (Levy and Razin 2019). Baccara and Yariv (2016) examine the conditions under which segregation results either in homogeneous groups or polarization. Our model assumes the segregation of society into echo chambers to be exogenous, and we shift our focus to the transmission of beliefs within existing clusters, and provide behavioral identification and characterization of this process.

There is a substantial body of recent literature that studies the process of social influence from a choice theoretic perspective (Fershtman and Segal 2018; Kashaev and Lazzati 2019; Lazzati 2020; Borah and Kops 2018; Chambers et al. 2019; Chambers et al. 2021; Cuhadaroglu 2017). In particular, the structure of our model draws inspiration from Fershtman and Segal (2018). Whereas we model influence through belief transmission, they look at influence in tastes. They consider two sets of preferences for each individual, represented by her core and behavioral utilities, and an influence function such that her behavioral utilities can be represented as a function of her core utility and others' behavioral utilities. Analogous to their model, core beliefs are private and behavioral beliefs are observable in our model.

Our work also connects with an extensive literature on social learning. Boundedly rational learning mechanisms have been argued to have both positive and negative effects on the accuracy of beliefs.[1] For instance, Glaeser and Sunstein (2009) outline situations in which such learning mechanisms propagate extremism in beliefs. Conversely, Martinez and Tenev (2024) argue that restricting information acquisition within echo chambers could actually improve the process of learning if the quality of various sources of information is uncertain. Likewise, Levy and Razin (2015a, 2015b) argue that correlation neglect does not necessarily increase polarisation in voting outcomes and might improve the accuracy of voting relative to a rational electorate. Acemoglu et al. (2014) characterize conditions under which *social cliques* can assist in asymptotic learning when costly communication networks are formed endogenously.

A common feature in the models detailed above is that private information or signals can be shared through the network. In contrast, a very popular strand of the literature details heuristics that may be used to learn from observed beliefs (DeGroot 1974; DeMarzo et al. 2003; Friedkin and Johnsen 1990; Levy and Razin 2018). Our model of behavioral beliefs and the underlying observational learning process are closely linked to such learning rules. Eyster and Rabin (2010) also model individuals who learn by sequentially observing the actions of those who immediately precede them. They find that when agents erroneously assume that the actions of their predecessors are based solely on private information, they may converge to incorrect actions with confidence. On the contrary, Bala and Goyal (1998) outline a rational process of local learning that may facilitate convergence to optimal actions under certain neighborhood structures. They find that the existence of highly influential nodes hinders accurate aggregation, similar to the conditions necessary for wise societies in the DeGroot model, as detailed by Golub and Jackson (2010).

The rest of the paper is organized as follows. The next section introduces our setup and formally defines an echo chamber representation. Section 3 provides the grounds for falsifiability of the model by characterizing it based on two conditions. Section 4 details the analysis of the identification of the model parameters. Section 5 presents the observational learning foundation of the model. Finally, in Section 6, we analyze the process of belief updating in our model and its application to information aggregation and voting. Proofs of all results are available in the Appendix.

[1] There is also a large class of papers that detail Bayesian models of social learning (e.g. Acemoglu et al. 2011, Jiménez-Martínez 2015, Azomahou and Opolot 2014). However, they tend to be both intractable and too demanding of observed behavior.

2 Setup

2.1 Primitives

Let S be a finite set of states, with any subset of S referred to as an event. Our stylized society consists of a set $I = \{1, \ldots, n\}$ of individuals. Each individual $i \in I$ is probabilistically sophisticated and her behavior over uncertain prospects is guided by a probability distribution π_i over the state space. We refer to this probability determining her behavior as her *behavioral beliefs*. We assume that an analyst or outside observer is able to observe these behavioral beliefs, say, from observing her betting behavior. Accordingly, the profile of behavioral beliefs $(\pi_i)_{i \in I}$ is data for the analyst.

2.2 Echo chamber representation

We now lay out our theoretical representation of interactions in the model that determine the behavioral beliefs of individuals. To that end, assume that the set of individuals in I are partitioned into the sets $\mathcal{E} = \langle E_1, \ldots, E_k \rangle$, with each element of the partition denoting an echo chamber (or chamber, for short) in society. To keep the setup meaningful, we assume that none of the echo chambers is a singleton. For any $i \in I$, we let $E(i)$ denote the echo chamber to which individual i belongs.

An individual's behavioral beliefs are formed both from her independent perception about the underlying uncertainty as well as the influence cast by the echo chamber she is a part of. Specifically, we assume that any such decision maker (DM), $i \in I$, is endowed with some *core beliefs* μ_i on S. Behavioral beliefs, of course, depend on core beliefs. But additionally, the working of influence within her echo chamber implies that her behavioral beliefs may be influenced by her perception of the overall beliefs prevailing in her echo chamber. We consider the average behavioral belief prevailing within her echo chamber, $\frac{1}{|E(i)|} \sum_{j \in E(i)} \pi_j$, as an aggregate summary statistic capturing this aspect of influence. One may think of this aggregate as the *echo* generated in the process of interactive belief formation.[2] We assume that the dependence on the two takes a linear weighted average form, with the weights determined by a parameter $\alpha_i \in (0, 1)$ that captures the degree to which this DM is immune to influence, i.e., higher is α_i, the less susceptible is this DM

[2]To understand better the idea behind the echo, suppose there is a change in the core beliefs of a DM i, say, due to fresh information she receives. This would clearly impact her behavioral beliefs and, in doing so, impact the behavioral beliefs of all others in her chamber, given the interactive nature of these beliefs. But, at the same time, the change induced in the behavioral beliefs of others would in turn bounce back like an echo and further influence i's behavioral beliefs, and so forth.

to influence. Specifically, for any event $A \subseteq S$, we assume that her behavioral belief is given by,

$$\overbrace{\pi_i(A)}^{\textbf{Behavioral belief}} = \alpha_i \underbrace{\mu_i(A)}_{\textbf{Core belief}} + \underbrace{(1-\alpha_i)\frac{1}{|E(i)|}\sum_{j \in E(i)} \pi_j(A)}_{\textbf{Echo}}$$

This, therefore, makes behavioral beliefs interactive within an echo chamber. Our representation of the profile of behavioral probabilities $(\pi_i)_{i \in I}$ provides a formal statement of these interactions. It captures the steady state of this process of interactions by requiring mutually consistent behavioral beliefs within an echo chamber.

Definition 1. *The profile of behavioral probability measures* $(\pi_i)_{i \in I}$ *on* S *has an* ***echo chamber representation*** *if there exists a partition* $\mathcal{E} = \langle E_1, E_2, ..., E_k \rangle$ *of* I*, and for each* $i \in I$*:*

- *a core probability measure* μ_i *on* S*, and*
- *an immunity from influence parameter* $\alpha_i \in (0, 1)$

such that $(\pi_i)_{i \in I}$ *can be defined as a solution to the system of equations,*

$$\pi_i(A) = \alpha_i \mu_i(A) + (1-\alpha_i)\frac{1}{|E(i)|}\sum_{j \in E(i)} \pi_j(A), \ \ i \in I,$$

Like in any equilibrium or steady state notion, we close the interactions by assuming that individuals hold correct expectations about the behavioral beliefs of others in their echo chamber. This allows them to correctly forecast the average behavioral belief about any event in their echo chamber.

We also focus on a subclass of echo chamber representations that identify the maximum extent of influence consistent with a given profile of behavioral beliefs. Observe that there are two channels that mediate the scope of influence in an echo chamber representation. The first is through the magnitude of the immunity of influence parameters $(\alpha_i)_{i \in I}$, the lower these are the greater is the scope of influence on individuals in an echo chamber. Accordingly, if there are two echo chamber representations with the same echo chamber partitioning and immunity to influence parameters $(\alpha_i)_{i \in I}$ and $(\tilde{\alpha}_i)_{i \in I}$, respectively, then we can say that the first representation doesn't fully capture the scope of influence consistent with the data if $\tilde{a}_i \leq \alpha_i$, for all i, holding strictly for some i. The second channel through which influence work is the size of echo chambers, the larger these are the more the number of individuals any given individual is influenced by and influences. Therefore, if under two echo chamber representations, the echo chamber partitioning in the second is a coarsening of the first, then we say that the first representation doesn't fully capture the scope of influence on this dimension. This motivates the following definition.

Definition 2. *An echo chamber representation $(\mathcal{E}, (\mu_i, \alpha_i)_{i\in I})$ of $(\pi_i)_{i\in I}$ is α-maximal if there does not exist another such representation $(\mathcal{E}, (\tilde{\mu}_i, \tilde{\alpha}_i)_{i\in I})$ of $(\pi_i)_{i\in I}$ s.t. $\tilde{\alpha}_i \leq \alpha_i$, for all $i \in I$, holding strictly for some i. An α-maximal representation $(\mathcal{E}, (\mu_i, \alpha_i)_{i\in I})$ is* ***maximal influence*** *if there does not exist another α-maximal representation $(\tilde{\mathcal{E}}, (\tilde{\mu}_i, \tilde{\alpha}_i)_{i\in I})$ s.t. $\tilde{\mathcal{E}}$ is a coarsening of $\mathcal{E}$.*

Remark 1 (Existence)**.** It is straightforward to establish that the steady state notion captured by the echo chamber representation doesn't suffer from concerns about non-existence. That is, given a collection $(\mu_i)_{i\in I}$ of core beliefs, it is immediate to establish that there exists a collection of behavioral beliefs $(\pi_i)_{i\in I}$ that simultaneously satisfy the system of equations:

$$\pi_i(A) = \alpha_i \mu_i(A) + (1-\alpha_i)\frac{1}{|E(i)|}\sum_{j\in E(i)} \pi_j(A), \ \ i \in I$$

To see this, note that the equation determining the behavioral beliefs of individual i, depends only on the beliefs of the individuals belonging to $E(i)$. Thus, it is sufficient to prove existence for a single echo chamber. Given a chamber E, subtract both sides by the average behavioral belief, and sum over all $i \in E$ to yield:

$$\frac{1}{|E|}\sum_{j\in E}\pi_j(A) = \sum_{i\in E}\frac{\alpha_i \mu_i(A)}{\sum_{j\in E}\alpha_j}$$

Substituting this expression in the earlier system of equations, we get:

$$\pi_i(A) = \alpha_i \mu_i(A) + (1-\alpha_i)\sum_{j\in E}\frac{\alpha_j \mu_j(A)}{\sum_{l\in E}\alpha_l}$$

This process can be used for each echo chamber, and the resulting collection $(\pi_i)_{i\in I}$ satisfies the system of equations. Since this is also true of all $(\pi_i)_{i\in I}$ that solve the equations, given a collection of $(\mu_i)_{i\in I}$, $(\alpha_i)_{i\in I}$ and partition $\mathcal{E} = \langle E_1, ..., E_k\rangle$, the resultant $(\pi_i)_{i\in I}$ must additionally be unique.

Remark 2 (Elite influence and α)**.** A key feature that our model demonstrates with respect to the immunity parameter, α, is that within each echo chamber, the individuals who are the least influenced (high α-s) happen to be the ones who end up having the greatest influence in terms of shaping beliefs within their echo chamber. This is in line with a theme that has featured prominently in recent times: elite influence. For instance, it has been pointed out in the context of partisan politics that each side of the partisan divide has elites who have a disproportionate influence on their respective sides. In other words, when there is influence at play, it is typically marked by a great degree of heterogeneity in terms of the ability to influence. Such an effect shows up in our model.

The nature of linear influence in the model gives the average behavioral belief in the echo chamber a unique structure. Recall the following expression for this average belief we derived in Remark 1:

$$\overline{\pi}_i(A) = \sum_{j \in E(i)} \frac{\alpha_j}{\sum_{k \in E(i)} \alpha_k} \mu_j(A)$$

That is, the average belief in an echo chamber can be represented as the weighted average of core beliefs, with the weights capturing relative influence. In particular, the weight attached to i's core belief is given by $\frac{\alpha_i}{\sum_{k \in E(i)} \alpha_k}$. It is relative because it depends on the ratio of the DM's own α_i to the sum of all α_j in her chamber. It also measures the degree of influence as the more a DM is immune to influence, the greater the weight placed on her core belief in the determination of the average behavioral belief. However, the more others are immune from influence, the more they influence the average belief, thus reducing the relative influence exhibited by the DM. A way of capturing the influence exhibited by a DM is the difference between her core beliefs and the average behavioral beliefs, which is expressed as follows.

$$|\overline{\pi}_i(A) - \mu_i(A)| = \left| \frac{1}{\sum_{k \in E(i)} \alpha_k} \left(\sum_{j \in E(i) \setminus \{i\}} \alpha_j (\mu_j(A) - \mu_i(A)) \right) \right|$$

This is a measure of her influence because it captures how close average behavioral beliefs in her echo chamber are pulled towards her core beliefs. Note that the difference is decreasing in α_i, and for any collection $(\mu_i)_{i \in I}$, the average behavioral belief is influenced more by i's core belief if she is more immune to influence.

3 Characterization

3.1 Maximal influence echo chamber representation

We now show that maximal influence echo chamber representations can be characterized by two axioms. The first of these axioms spells out the key constraint that influence imposes on a DM's behavioral beliefs. It says that a DM is constrained by influence to think of an event as sure only if there exists at least some other individual who does likewise, with this influence being mutual. In other words, even if an individual fundamentally thinks of an event as sure by her core beliefs, it is not guaranteed to translate into her behavior unless it receives conformity from others, with this process of conformity being mutual.

Axiom A1 (Certainty conformism)**.** *For all $i \in I$, there exists $j \in I$, $j \neq i$, such that, for any event $A \subseteq S$, $\pi_i(A) = 1$ if and only if $\pi_j(A) = 1$.*

The first axiom incorporates the idea of conformism, which the literature identifies as a key marker of social influence. When specialized to our current context, such conformism is seen on behavioral probabilistic judgments about sure events.

Our second axiom, on the other hand, reinforces the message that inter-personal interactions of influence need not be just about conformity. There is scope for any individual influencing the beliefs of others she interacts with and not just mimicking them. To present this axiom, we first introduce a definition. We define the revealed neighborhood of any $i \in I$ by

$$N(i) = \{j \in I \setminus \{i\} : \pi_i(A) = 0 \iff \pi_j(A) = 0, \text{for any event } A\}$$

The idea behind this revealed elicitation is quite straightforward. An individual is presumably connected to those individuals who she seeks conformity from. Hence, these individuals form her neighborhood in the social network of interacting beliefs. Observe that, the certainty conformism axiom guarantees that $N(i)$ is non-empty for any i. Now for any $J \subseteq I$, and any event A, denote the average behavioral belief about A amongst individuals in J by

$$\overline{\pi}_J(A) = \frac{1}{|J|} \sum_{j \in J} \pi_j(A)$$

Axiom A2 (Subjective assessment of uncertainty)**.** *For all $i \in I$, there exists an event A such that $\overline{\pi}_{N(i)}(A) \neq \overline{\pi}_{N(i)\cup\{i\}}(A)$*

The axiom highlights the point that any $i \in I$ is not simply a passive recipient of the probabilistic judgments of her neighbors. Rather, her presence may also change the average beliefs prevailing in her neighborhood in as much as there exists at least some event in which the average belief in her neighborhood with and without her is different. Equivalently, on this event, i's behavioral belief is different from the average behavioral belief in her neighborhood, i.e., $\pi_i(A) \neq \overline{\pi}_{N(i)}(A)$. Observe that, given what the first axiom says, such influence that i casts is necessarily for events that she and everyone in her neighborhood share a degree of uncertainty over. The axiom can also be seen in the light of "maximality" of influence. It implies that all individuals cast some influence on the formation of beliefs. There is no individual who is simply the recipient of influence.

We can now present our behavioral characterization result.

Theorem 1. *The behavioral probability profile $(\pi_i)_{i\in I}$ has a maximal influence echo chamber representation if and only if it satisfies certainty conformism and subjective assessment* *of uncertainty.*

Proof: Please refer to Appendix Section A.4

3.2 Inter (Intra) chamber restrictions on core beliefs

The primitives of our model do not impose any interpersonal restrictions on the core beliefs of individuals, neither within an echo chamber nor across chambers. In reality, one would imagine that these beliefs may be inter-related to some extent. Presumably, a necessary condition for individuals coming together and forming an echo chamber is some minimal agreement in their core beliefs, especially when it comes to perceptions about certainty, with this agreement not shared by those outside the chamber. This is in the spirit of belief-based homophily. At the same time, for these echo chambers to retain their salience, pathways of influence must exist within them and it is this influence that serves as a glue holding them together. But for influence to operate, there also needs to be disagreements within the echo chamber that opens up the scope for such influence. In other words, we would imagine that the functioning of echo chambers incorporate both inter-chamber and intra-chamber disagreements over core beliefs. It is, therefore, interesting to note that one of the things that the characterization of the maximal influence echo chamber model brings to the forefront is that such notions of disagreements are embedded in this class of representations.

To state these observations formally, we introduce some terminology. We say that i is fundamentally certain about an event A if $\mu_i(A) = 1$. We say that a chamber E is fundamentally certain about an event A if $\mu_i(A) = 1$, for all $i \in E$. We now introduce two conditions pertaining to disagreements on fundamental certainty (DFC). The first condition provides a statement about differing views on certainty across echo chambers.

Condition 1 (DFC-Inter). *For all chambers E and E', there exists an event that one of the chambers is fundamentally certain about, but the other chamber is not.*

As mentioned above, one may think of this as a constitutive condition of echo chambers—the fact that, at a fundamental level, echo chambers are formed by people who share a certain view of reality that is not necessarily shared by those outside the echo chamber. The condition above expresses this idea in a fairly weak form. At the same time, as noted earlier, just because echo chambers form presumably doesn't mean that there is complete agreement on matters of certainty within an echo chamber. Our next condition captures this viewpoint, leaving open the possibility that there is room for disagreement and influence within an echo chamber.

Condition 2 (DFC-Intra)**.** *For all $i \in I$, there exists an event she is fundamentally certain about, but someone in her chamber is not*

Theorem 2. *Suppose $(\mathcal{E} = \langle E_1, \ldots, E_k \rangle, (\mu_i, \alpha_i)_{i \in I})$ is an echo chamber representation of $(\pi_i)_{i \in I}$. $(\mathcal{E} = \langle E_1, \ldots, E_k \rangle, (\mu_i, \alpha_i)_{i \in I})$ is a maximal influence echo chamber representation iff it satisfies DFC-Inter and DFC-Intra.*

Proof: Please refer to Appendix Section A.6

The result establishes that DFC-Inter and DFC-Intra are properties of maximal influence echo chamber representations and not of echo chamber representations in general. This adds to the appeal of this sub-class of representations as they implicitly capture intuitive notions underling echo chamber formation and salience, even though we do not explicitly model this in our set-up.

3.3 Echo chamber representation

A final detail regarding characterization that the reader may be interested in is about the characterization of echo chamber representations. It turns out such representations are characterized simply by certainty conformism.

Theorem 3. *The behavioral probability profile $(\pi_i)_{i \in I}$ has an echo chamber representation if and only if it satisfies certainty conformism.*

Proof: Please refer to Appendix Section A.3

4 Identification

We now address the question about the identification of the model parameters. That is, suppose $(\pi_i)_{i \in I}$ has a maximal influence echo chamber representation. Is this representation unique, or are multiple such representations possible? The following result shows that a desirable feature of such a representation is that it is uniquely identified.

Theorem 4. *If $(\mathcal{E} = \langle E_1, \ldots, E_k \rangle, (\mu_i, \alpha_i)_{i \in I})$ and $(\tilde{\mathcal{E}} = \langle \tilde{E}_1, \ldots, \tilde{E}_l \rangle, (\tilde{\mu}_i, \tilde{\alpha}_i)_{i \in I})$ are maximal influence echo chamber representations of $(\pi_i)_{i \in I}$, then $\mathcal{E} = \tilde{\mathcal{E}}$, and for each $i \in I$, $\mu_i = \tilde{\mu}_i$, and $\alpha_i = \tilde{\alpha}_i$.*

Proof: Please refer to Appendix Section A.5

The exact identification of maximal influence echo chamber representations contrasts sharply with echo chamber representations where we don't impose this restriction. Such representations may not be precisely identified and the best we are able to do is provide bounds on the echo chamber partitioning and immunity to influence parameters.

We first examine the possible partitionings of society that can be accommodated under an echo chamber representations of a profile of behavioral probabilities $(\pi_i)_{i\in I}$. We show that if $\mathcal{E}$ is such a partitioning then for any $j \in E(i)$, $j \neq i$, $j \in N(i) = \{j \in I \setminus \{i\} : \pi_i(A) = 0 \iff \pi_j(A) = 0, \text{for any event } A\}$. That is, $E(i)$ must be a subset of $N(i) \cup \{i\}$ for all $i \in I$. Further, $N(i) \cup \{i\}$ is the largest possible echo chamber to which i may belong.

Proposition 1. *If $\tilde{\mathcal{E}} = \left\langle \tilde{E}_1, ..., \tilde{E}_\ell \right\rangle$ is a partition of society under an echo chamber representation of the beliefs $(\pi_i)_{i\in I}$, then $\tilde{E}(i) \subseteq N(i) \cup \{i\}$. Further, if $\tilde{E}(i) \subseteq N(i) \cup \{i\}$ and $|\tilde{E}(i)| \geq 2$ for all $i \in I$, then there exists an echo chamber representation of $(\pi_i)_{i\in I}$ with the partition $\tilde{\mathcal{E}}$.*

Proof: Please refer to Appendix Section A.1

Next, we identify the range for the immunity from influence parameters consistent with an echo chamber representation of behavioral beliefs and establish a lower bound on these.

Proposition 2. *Suppose the collection of behavioral beliefs $(\pi_i)_{i\in I}$ has an echo chamber representation with the partition $\mathcal{E} = \langle E_1, ..., E_\ell \rangle$, and $E(i)$ denoting i's echo chamber under $\mathcal{E}$. Then $(\pi_i)_{i\in I}$ has an echo chamber representation with influence parameters $(\alpha_i)_{i\in I}$ iff $\alpha_i \in \left[1 - \min \frac{\pi_i(A)}{\bar{\pi}_{E(i)}(A)}, 1\right) \cap (0, 1)$ for all $i \in I$.*

Proof: Please refer to Appendix Section A.2

An insight we can gain from this result is that once we fix a partition $\mathcal{E}$, restrictions on the influence parameters of any $i \in I$ depend only on $E(i)$. However, while the bounds on these parameters are dependent on $E(i)$, the exact choice is independent of all other individuals. Particularly, what this means is that if, under some partition $\mathcal{E}$, $(\alpha_i)_{i\in I}$ and $(\tilde{\alpha}_i)_{i\in I}$ are two collections of influence parameters that can be accommodated under an echo chamber representations of $(\pi_i)_{i\in I}$, then so can be $(\alpha_i, \tilde{\alpha}_{-i})$ and $(\tilde{\alpha}_i, \alpha_{-i})$, for any $i \in I$.

5 Behavioral probabilities from observational learning

The next task we undertake is to connect our theory of behavioral probabilities to the observational learning literature that follows from DeGroot (1974). Indeed, we provide an observational learning foundation for the echo chamber model as presented in the earlier sections, and for a theory of belief updating under it that we consider subsequently. The formulation of observational learning that we propose combines the basic insight of DeGroot learning with an environment where agents receive private information in a sequential setting. That is, unlike DeGroot, where agents receive all relevant private information as captured by their initial beliefs before any communication take place between them, we consider an environment where interactions and the arrival of private information overlap.

In the DeGroot model, individuals start with some initial beliefs, π^0, and update their beliefs in every period by taking a weighted average of the last observed beliefs of all individuals. Let W_{ij} denote the weight that individual i places on the beliefs of individual j. $W_{ij} \geq 0$ for all i, $j \in I$, and $\sum_j W_{ij} = 1$ for all $i \in I$. Individual i is connected to j if $W_{ij} > 0$. These weights remain constant over time, so the updating process can be written as

$$\pi^{t+1}(A) = W\pi^t(A)$$

where $\pi^t(A)$ and $\pi^{t+1}(A)$ are vectors of the beliefs of all individuals about event A in periods t and $t+1$, and W is the $n \times n$ matrix of weights.[3] In a network where the subnetwork on each component is strongly connected (true of echo chambers), this model would imply consensus amongst the individuals of each component (Golub and Jackson 2010). That is $\lim_{t\to\infty} \pi_i^t(A) = \lim_{t\to\infty} \pi_j^t(A)$ for all $A \subseteq S$ if i and j belong to the same component. We modify this by introducing a component of beliefs purely based on private information, which is subject to being updated as new information arrives.

In our model, individuals hold distinct private and reported beliefs, which evolve according to the process

$$\pi^{t+1}(A) = \Lambda\mu^{t+1}(A) + (I_n - \Lambda)W\pi^t(A)$$

Here, Λ is a diagonal matrix with $\Lambda_{ii} \equiv \lambda_i$, which denotes the weight that individual i places on her private belief. I_n is the $n \times n$ identity matrix, so she places $1 - \lambda_i$ weight on the DeGroot updated belief. $\mu^{t+1}(A)$ is the vector of private beliefs about event A, which is subject to being updated based on new information that may arrive in each period.

[3]Note, in general, the DeGroot model is used to capture the evolution of scalar beliefs, typically estimates of parameters, and the convergence of these estimates to a particular value through the updating process.

Suppose that private information is conveyed to i in each period $t < T_i$ in the form of a tuple (C, ϕ_i) for each $i \in I$. C is a finite set of signals and $\phi_i : S \to \Delta C$ is a mapping that defines the probability, $\phi_i^s(c)$, of individual i observing the signal $c \in C$, conditional on the true state being s. A single signal is perfectly informative if ϕ_i is an injective mapping and maps to degenerate distributions in ΔC. Signals are completely uninformative if ϕ_i is a constant mapping. Individuals hold some initial private beliefs $(\mu_i^0)_{i\in I}$ and update these beliefs every period upon receiving new information according to Bayes rule.[4] We maintain that signals cannot be conveyed to anyone else, which means that every individual's private beliefs are updated using only the signals they receive. Any communication of information occurs purely through reported beliefs. Furthermore, if no new information is conveyed to an individual, her private beliefs remain unchanged.

Denote individual i as having a link to j if $W_{ij} > 0$. Since the beliefs of individuals depend only on the beliefs of individuals to whom they have a path, we can reduce our analysis to a single component, E, at once.[5] The updating process is then reduced to

$$\pi_E^{t+1}(A) = \Lambda_E \mu_E^{t+1}(A) + (I_{|E|} - \Lambda_E) W_E \pi_E^t(A)$$

Since we are interested in cluster networks, we analyze a particular restriction of this model with an undirected, unweighted network, in which individuals in the same component always have a direct link to each other. That is, the subnetwork on each component is complete. Let $\mathcal{E} = \langle E_1, ..., E_k \rangle$ denote the partition of I into components. Supposing that individuals account for their own beliefs as well in the component average, the adjacency matrix G can be represented with $G_{ii} = 1$, and the weighting matrix W is the degree-adjusted adjacency matrix.[6] The weighting matrix, W_E, for the component E is then $|E|^{-1}\mathbf{1}_{|E|}$,[7] and an individual's reported beliefs are updated every period according to

$$\pi_i^{t+1}(A) = \lambda_i \mu_i^{t+1}(A) + (1 - \lambda_i) \frac{1}{\sum_{j\in I} G_{ij}} \sum_{j\in I;\, G_{ij}=1} \pi_j^t(A)$$

Since the stream of private information is finite, the sequence of private beliefs $\{\mu^t(A)\}_{t=0}^{\infty}$ is eventually constant for any possible realization of signals and each $A \subseteq S$. Building on this fact, we show that $\{\pi^t(A)\}_{t=0}^{\infty}$ converges to some π for all initial reported beliefs π^0. We then establish that π is a profile of behavioral beliefs that admit an echo chamber

[4] Given c^i as the signal received by i, her updated private belief is given by:

$$\mu_i(A|c^i) = \frac{\sum_{s\in A} \mu_i(s) \phi_i^s(c^i)}{\sum_{s'\in S} \mu_i(s') \phi_i^{s'}(c^i)}$$

[5] Individuals i and j are in the same component if at least one of them has a path to the other.

[6] G is the adjacency matrix of the network if $G_{ij} = 1$ whenever there exists a link between i and j. Each row in W is obtained by scaling the corresponding row in G such that the row sums to one.

[7] $\mathbf{1}_n$ is the $n \times n$ matrix of ones.

representation.

Proposition 3. *Suppose individuals in network $\mathcal{G}$ are divided into complete components $\mathcal{E} = \langle E_1, ..., E_k \rangle$. Given initial private beliefs $(\mu_i^0)_{i\in I}$ and reported beliefs $(\pi_i^0)_{i\in I}$, define reported beliefs for each $i \in I$ in period $t+1$, $t \in \mathbb{Z}_+$, as*

$$\pi_i^{t+1}(A) \equiv \lambda_i \mu_i^{t+1}(A) + (1-\lambda_i)\frac{1}{\sum_{j\in I} G_{ij}} \sum_{j\in I;\, G_{ij}=1} \pi_j^t(A)$$

Then, for each $i \in I$, $\{\pi_i^t\}_{t=0}^{\infty}$ is convergent with $\lim_{t\to\infty} \pi_i^t = \pi_i$. Furthermore, $(\pi_i)_{i\in I}$ is the profile of behavioral beliefs that admits an echo chamber representation with parameters $\mathcal{E} = \langle E_1, ..., E_k \rangle$, $(\mu_i)_{i\in I}$ such that $\mu_i = \mu_i^{T_i}$, and $(\alpha_i)_{i\in I}$ with $\alpha_i = \lambda_i$ for all $i \in I$.

Proof: Please refer to Appendix Section A.7

Propositon 3 shows that behavioral beliefs in the Echo Chamber model can be thought of as the limit of a boundedly rational learning process as described above. This means that the axioms we provide to characterize echo chamber representations can be used to analyze the parameters of such an updating process even with data from only the limiting distribution.

The result also highlights two important features that connect our model with findings in the observational learning literature. Firstly, our model can be thought of as a special case of the model by Friedkin and Johnsen (1990). Once all new information has arrived, private beliefs remain constant, equivalent to the anchored initial beliefs in their model. The weighting matrices applied to the anchored initial beliefs and previous-period beliefs are given by Λ and $(I_n - \Lambda)W$. The weighting matrices in our model satisfy their conditions for convergence by definition. However, we relax their stipulation of initial reported beliefs being equal to the weighted private beliefs in our convergence result.

The second detail is that of the importance of information sequencing and the convergence of behavioral beliefs. Reshidi (2024) shows in a model with one-time information release for each individual and naive incorporation of private information, that the sequence in which information is conveyed affects the limiting beliefs from a subsequent DeGroot process. In Reshidi's model, each individual receives a signal s_i in some period t_i. They then update their beliefs according to the rule

$$\pi_i^{t+1} = \begin{cases} \lambda_i(H^t)s_i + (1-\lambda_i)(H^t)W_i\pi^t & t = t_i - 1 \\ W_i\pi^t & \text{otherwise} \end{cases}$$

where H^t is the history of signals released up until period t, W_i is the i-th row of the

DeGroot weighting matrix, and $\lambda_i(H^t)$ is a history-dependent weight placed on new private information. When individuals do not receive any new information, they follow the DeGroot method of updating their beliefs. On the other hand, if they receive new information, they take a weighted average of the signal and the DeGroot updated belief. Reshidi shows that there do not exist any weights for which the limiting belief from this process is sequence independent.

The intuition behind this result, and the contrast with our model, lies in how private information is incorporated and how beliefs are updated thereafter. Particularly, Reshidi's model is identical to that of DeGroot once the last individual receives new information. In the simple case of an aperiodic, strongly connected network, we know that the DeGroot model leads to convergence of beliefs to a weighted average of initial beliefs. However, even in the simple case of history-independent weights, the sequence of information release before period $t_i - 1$ determines π^{t_i-1}, which determine "initial" beliefs of the subsequent DeGroot process, if i is the last individual who receives new information.

In our setup, however, new information is incorporated into private beliefs in a Bayesian manner, without being affected by others' reported beliefs, which makes them sequence-independent. Since updated reported beliefs incorporate private beliefs in every period, the limit of reported beliefs is dependent solely on final private beliefs. Therefore, the limiting reported beliefs are also independent of the sequencing of private information.

6 Belief Updating and Information Aggregation

6.1 Belief Updating and Transmission

We now examine how new information is assimilated and transmitted in our model. Suppose that private information is conveyed once to individuals in the manner described in the previous section. We examine how this information may be incorporated into individuals' beliefs, how these beliefs are transmitted to others in the echo chamber, and the behavioral implications of this mechanism. We first show that beliefs are sticky, and respond fully to nothing but consistent information across the echo chamber. We then relate this to confirmation bias and its effects on voting outcomes.

Let $\mathcal{S}$ denote the set of non-empty subsets of the state space S. For each $i \in I$, define $\sigma_i : C \to \mathcal{S}$ such that $\sigma_i(c) = \{s \in S : \phi_i^s(c) > 0, \, \mu_i(s) > 0\}$. For all $c \in C$, we assume that $\sigma_i(c) \neq \emptyset$. That is, $\phi_i^s(c) > 0$ for some $s \in S$ such that $\mu_i(s) > 0$.

First, consider the case when a single individual, i, receives new information, c^i. Let π'_j denote individual j's posterior behavioral belief upon receiving the signal. By Remark 1, the individual's revised belief is given by

$$\pi'_i(A) = \pi_i(A) + \alpha_i(\mu_i(A|c^i) - \mu_i(A)) + (1-\alpha_i)\frac{\alpha_i(\mu_i(A|c^i) - \mu_i(A))}{\sum_{j\in E(i)} \alpha_j}$$

As is apparent from the expression, individuals' responsiveness to new information is dampened by their susceptibility to influence and bolstered by their relative influence in the echo chamber. Note also from the above expression that $\frac{\pi'_i(A)-\pi_i(A)}{\mu'_i(A|c^i)-\mu_i(A)} = \alpha_i + \frac{\alpha_i(1-\alpha_i)}{\sum_{k\in E(i)}\alpha_k}$. As such, the effect of α on an individual's responsiveness to new information, relative to core beliefs, can be decomposed into two components: Direct effect, α_i, and Echo effect, $\frac{\alpha_i(1-\alpha_i)}{\sum_{k\in E(i)}\alpha_k}$. This is illustrated in Figure 1.

Figure 1. Impact of α on new information: Direct and Echo effect

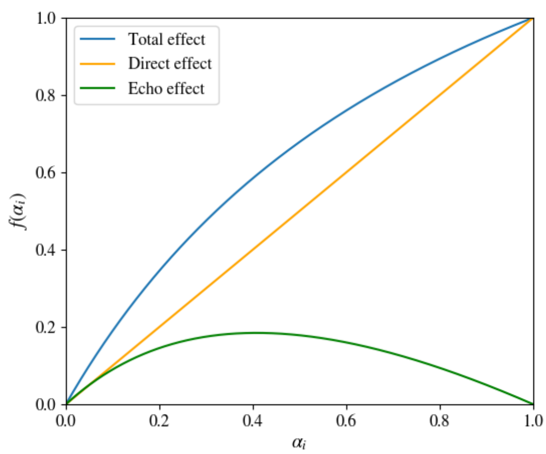

Far from the Bayesian benchmark, however, not only is information not fully incorporated, behavioral beliefs might actually be updated contrary to private information when others in the echo chamber receive opposing information. Particularly, suppose $A \cap \sigma_i(c^i) = \emptyset$. A Bayesian DM would revise their belief to $\mu_i(A|c^i) = 0$. However, individuals in echo chambers are unable to incorporate such information fully, and might actually revise their beliefs upward. In fact, $\pi'_i(A) > \pi_i(A)$ for such A if

$$\frac{\alpha_i(1+\sum_{j\in E(i)\setminus\{i\}}\alpha_j)}{1-\alpha_i}\mu_i(A) < \sum_{j\in E(i)\setminus\{i\}} \alpha_j(\mu_j(A|c^j) - \mu_j(A))$$

This particular inequality suggests that individuals especially struggle to rule out the possibility of low probability states upon receiving information, particularly because they seek consensus on impossible events within their echo chamber. Therefore, whenever others receive conflicting information, low (immunity to) influence individuals are susceptible to ignoring their private information.

We would thus want to identify the restrictions on the information source that we must place if individuals are to update akin to Bayesians, at least on sure (therefore, null) events. The result that follows shows that in an echo chamber setting, even the slightest heterogeneity in privately conveyed information leads to a conflict between an individual's information and the observed beliefs of other individuals in the chamber.

Proposition 4 (Sticky Beliefs)**.** *For any $i \in I$, $\pi_i'(\sigma_i(c^i)) = 1$ if and only if $\mu_j(\sigma_i(c^i) \mid c^j) = 1$, for all $j \in E(i)$. Then $\pi_i'(\sigma_i(c^i)) = 1$ if and only if $\sigma_i(c^i) = \sigma_j(c^j)$ for all $j \in E(i)$.*

Proof: Please refer to Appendix Section A.8

The first part of the result is derived from the property of the model, that $\pi_i(A) = 1 \iff \mu_j(A) = 1$ for all $j \in E(i)$. The second result then follows from the first.

6.2 Information Aggregation and Voting

Keeping in mind the stickiness of beliefs and the occasional negligible effect of receiving private information, we analyze how information is aggregated in the echo chamber model. We do so in the context of a simple model of majority voting with common interests but differing beliefs between opposed groups of individuals. We show that influence in beliefs increases susceptibility to confirmation bias, which is exhibited in the form of a *perseverance effect.* We also show, in similar spirit to jury theorems,[8] that as the society grows arbitrarily large, the probability of the majority voting correctly converges to 1 if the source of information is sufficiently accurate. However, while the wisdom of the crowd prevails even with naïve updating of behavioral beliefs, it requires a greater accuracy of private information than in the absence of influence.

Suppose the society must choose between two policy regimes, L and R. Individuals, $I = \{1, ..., 2n\}$, are evenly divided into two chambers, E_L and E_R. Policy L can be thought of as the in-group policy for E_L, and likewise with policy R for E_R. Each individual votes for a policy, and one of the two policies is implemented by simple majority. Ties are broken by a coin toss.

There are two possible states of the world, $S = \{\ell, r\}$. If the true state of the world is ℓ and policy L is implemented, then all individuals receive $\overline{u}$, irrespective of the chamber to

[8] The canonical jury theorem is by Condorcet (1785), where votes are randomly and independently cast with uniform probability. Extensions of the Condorcet's theorem include contexts with correlated votes and strategic voting: e.g. Kaniovski and Zaigraev (2011), Dietrich (2008), Peleg and Zamir (2012).

which they belong. Likewise, if policy R is implemented and the true state is r. If L and R are implemented with the underlying true state being r and ℓ respectively, individuals receive the utility $\underline{u}$. Assume $\overline{u} > \underline{u}$.

Suppose that individuals have identical core beliefs within echo chambers. That is, if $i \in L$, $\mu_i = \mu_L$, and $\mu_i = \mu_R$ otherwise. Consequently, $\pi_i = \bar{\pi}_{E(i)} = \mu_{E(i)}$, where $E(i) \in \{L, R\}$. For simplicity, assume $\mu_L(\ell) = \mu_R(r) = q \in (\frac{1}{2}, 1)$. For simplicity, also assume that $\alpha_i = \alpha$ for all $i \in I$. Individuals receive conditionally independent signals, $C = \{c_\ell, c_r\}$, through a common source (C, ϕ) of accuracy $\rho \in [\frac{1}{2}, 1)$. That is, $\phi^\ell(c_\ell) = \phi^r(c_r) = \rho$. Upon receiving a signal, individual i casts a vote in favor of L if $\pi_i'(L) > \frac{1}{2}$, and likewise for R.[9] If $\pi_i'(\ell) = \pi_i'(r)$, she votes for her in-group policy.

Remark 3. *It is useful to establish the conditions under which individuals vote according to the private information they receive, in the absence of any influence. Particularly, when does the private signal observed by an individual fully determine her vote if she votes purely based on her core beliefs? The answer to this: when the accuracy (ρ) of the signal is greater than the strength of the prior in-group belief (q). To see this,*

$$\begin{aligned} \mu_L(r|c_r) &= \frac{\rho(1-q)}{\rho(1-q)+(1-\rho)q} > \frac{1}{2} \\ \iff 2\rho(1-q) &> \rho(1-q)+(1-\rho)q \\ \iff \rho(1-q) &> (1-\rho)q \\ \iff \rho &> q \end{aligned}$$

This provides a Bayesian benchmark against which we can compare the results from the Echo Chamber model.

We now analyze how voting based on behavioral beliefs is susceptible to confirmation bias. To do so, we first make an observation about how behavioral beliefs respond to signals received by the echo chamber. Since private information is not incorporated fully with influence in beliefs, an individual's vote is not only dependent on the signal she receives, but also on the distribution of signals observed by her echo chamber. Lemma 1 shows that individuals vote for the out-group policy if and only if a certain proportion of individuals in her echo chamber receive the out-group signal. When this threshold is met, all such individuals vote for the out-group policy together. However, when it is not, their in-group beliefs persevere despite evidence to the contrary.

Lemma 1. *Let X_ℓ^L denote the number of individuals in E_L who received the signal c_ℓ. If $\alpha \leq \frac{1}{2}$ and $\rho \neq 1$, there exists $\hat{x}(\alpha, \rho, q) \in [0, 1)$ such that conditional on receiving the*

[9]See that if this were modelled as a game, then voting for the policy they find is more likely to give them higher utility is a weakly dominant strategy.

signal c_r, $\pi_i'(\ell) < \frac{1}{2}$ *iff* $\frac{X_\ell^L}{n} < \hat{x}(\alpha,\, \rho,\, q)$ *for* $i \in E_L$.

Proof: Please refer to Appendix Section A.9

The result shows that the updated behavioral belief, $\pi_i'(\ell)$ on the in-group state is less than half, enough to induce a switch in i's vote, only when the proportion of individuals who receive the signal c_ℓ is less than $\hat{x}$, a quantity dependent on α, ρ, and q. While the above result is written for an individual in E_L, by the symmetry of the problem, the same holds true for out-group voting in E_R. The result that follows shows that this need not hold true when individuals are relatively immune from influence. They can sufficiently and independently incorporate private information that originates from accurate but imperfect signals. Define $\rho^*(\alpha,\, q) \equiv \min \operatorname{argmax}_\rho \hat{x}(\alpha,\, \rho,\, q)$ as the smallest value of ρ that maximizes $\hat{x}(\alpha,\, \rho,\, q)$ for a given α and q. The proof for Lemma 1 shows that $\hat{x}$ is continuous and increasing in ρ, which implies that $\rho^*(\alpha,\, q) = 1$ whenever $\alpha \leq \frac{1}{2}$. Lemma 2 shows that this does not hold when $\alpha > \frac{1}{2}$.

Lemma 2. *If* $\alpha > \frac{1}{2}$, $\rho^*(\alpha,\, q) \in (q,\, 1)$ *and* $\hat{x}(\alpha,\, \rho^*(\alpha,\, q),\, q) = 1$.

Proof: Please refer to Appendix Section A.10

Using this insight, we can ask how likely is it that a sub-optimal policy gets implemented when beliefs are influenced in echo chambers. Naturally, a policy is sub-optimal if it gives individuals utility $\underline{u}$, which, in turn, is state-contingent. With random realizations of signals, it is clear that even when $\rho > q$, voting under core beliefs may result in a positive probability of implementing the sub-optimal policy. Proposition 5 shows that unless voting behavior (in line with the earlier lemmas) is uninfluenced by echo chambers, the probability of implementing a sub-optimal policy is greater with influence than without. Let $\gamma(\alpha,\, \rho,\, q)$ and $\tilde{\gamma}(\alpha,\, \rho,\, q)$ denote the probability of implementing a sub-optimal policy under voting by behavioral and core beliefs, respectively, for a given α, ρ, and q.

Proposition 5. *Suppose* $q < \rho < \rho^*(\alpha,\, q)$ *and* n *finite but large. Then,* $\gamma(\alpha,\, \rho,\, q) > \tilde{\gamma}(\alpha,\, \rho,\, q)$. *Furthermore,* $\gamma(\alpha,\, \rho,\, q) \to \tilde{\gamma}(\alpha,\, \rho,\, q)$ *as* $\hat{x}(\alpha,\, \rho,\, q) \to 0$ *or* $\hat{x}(\alpha,\, \rho,\, q) \to 1$.

Proof: Please refer to Appendix Section A.11

The next question we are interested in is whether majority voting by behavioral beliefs benefits from the wisdom of the crowd. If so, what are the conditions under which this occurs? For reference, note that when $\rho > q$, individuals vote correctly with independent probability $\rho > q > \frac{1}{2}$. Then, by the Condorcet Jury Theorem, as $n \to \infty$, majority voting almost surely results in the correct policy being implemented. On the other hand, if $\rho \leq q$, all individuals simply vote for their in-group policy, in which case the implementation of the policy is always decided by coin-toss, leading to a $\frac{1}{2}$ probability of error. The next proposition shows that majority voting by behavioral beliefs also exhibits similar asymptotic characteristics, in that it is either almost surely correct or akin to a coin toss. Nevertheless, it demands greater degree of signal accuracy to effectively aggregate information.

Proposition 6. *For each α, q, there exists $\tilde{\rho}(\alpha,\, q) \in (q,\, 1)$ such that, as $n \to \infty$, if:*

1. $\rho \leq \tilde{\rho}(\alpha,\, q)$*, then* $\gamma(\alpha,\, \rho,\, q) \to \frac{1}{2}$
2. $\rho > \tilde{\rho}(\alpha,\, q)$*, then* $\gamma(\alpha,\, \rho,\, q) \to 0$

A Appendix

The proofs for all the results in the paper are contained here. We first prove propositions 1 and 2, as they are later used in the proofs of theorems 1, 2, and 4. We then prove theorem 3, followed by theorem 1, 4, and then theorem 2.

A.1 Proof of Proposition 1

Suppose $\left(\tilde{\mathcal{E}},\, (\tilde{\alpha}_i)_{i\in I},\, (\tilde{\mu}_i)_{i\in I}\right)$ is an echo chamber representation of $(\pi_i)_{i\in I}$. Note

$$\pi_i(A) = \tilde{\alpha}_i \tilde{\mu}_i(A) + (1-\tilde{\alpha}_i)\frac{1}{|\tilde{E}(i)|}\sum_{j\in\tilde{E}(i)}\pi_j(A)$$

By the fact that $\pi_j(A) \in [0,\, 1]$ and $\mu_i(A) \in [0,\, 1]$ for all A and j, through $\tilde{\alpha}_i \in (0,\, 1)$, we get $\pi_i(A) = 0$ iff $\pi_j(A) = 0$ for all $j \in \tilde{E}(i)$. This applies to all $i \in I$. Then if $j \in \tilde{E}(i)$, then it must be that $\pi_i(A) = 0 \iff \pi_j(A) = 0$, which means $j \in E(i)$. Thus $\tilde{E}(i) \subseteq E(i)$.

For the converse, suppose $\tilde{E}(i) \subseteq E(i)$ and $|\tilde{E}(i)| \geq 2$ for all $i \in I$. Let $\tilde{\bar{\pi}}_i(A) = \bar{\pi}_{\tilde{E}(i)}(A)$ for notational simplicity. By $\tilde{E}(i) \subseteq E(i)$, we have that $\pi_i(A) = 0$ iff $\pi_j(A) = 0$ for all $j \in \tilde{E}(i)$ iff $\tilde{\bar{\pi}}_i(A) = 0$. Then $\frac{\pi_i(A)}{\tilde{\bar{\pi}}_i(A)} > 0$ for all A such that this ratio is defined, which means that $1 - \min \frac{\pi_i(A)}{\tilde{\bar{\pi}}_i(A)} < 1$. Now choose some $\alpha_i \in \left(1 - \min \frac{\pi_i(A)}{\tilde{\bar{\pi}}_i(A)},\, 1\right)$. Since $\pi_i(S) = \tilde{\bar{\pi}}_i(S)$, we have $\min \frac{\pi_i(A)}{\tilde{\bar{\pi}}_i(A)} \leq 1$, which means $\alpha_i \in (0,\, 1)$.

Defining μ_i as:

$$\mu_i(A) = \frac{\pi_i(A) - (1-\alpha_i)\tilde{\bar{\pi}}_i(A)}{\alpha_i}$$

For A such that $\pi_i(A) = 1$ or $\pi_i(A) = 0$, we have $\mu_i(A) = 1$ and $\mu_i(A) = 0$ respectively. Further, our choice of α_i is such that $\min \frac{\pi_i(A)}{\tilde{\overline{\pi}}_i(A)} > 1 - \alpha_i$, which means $\pi_i(A) - (1 - \alpha_i)\tilde{\overline{\pi}}_i(A) > 0$ for all other A, implying that $\mu_i(A) \geq 0$ for all $A \subseteq S$. For A, B disjoint

$$\begin{aligned}
\mu_i(A \cup B) &= \frac{\pi_i(A \cup B) + (1 - \alpha_i)\overline{\pi}_i(A \cup B)}{\alpha_i} \\
&= \frac{\pi_i(A) + \pi_i(B) - (1 - \alpha_i)(\overline{\pi}_i(A) + \overline{\pi}_i(B))}{\alpha_i} \\
&= \frac{\pi_i(A) - (1 - \alpha_i)\overline{\pi}_i(A)}{\alpha_i} + \frac{\pi_i(B) - (1 - \alpha_i)\overline{\pi}_i(B)}{\alpha_i} \\
&= \mu_i(A) + \mu_i(B)
\end{aligned}$$

Thus, μ_i is a probability. Now that we have established that α_i and μ_i are valid choices for the parameters, rearranging the definition for $\mu_i(A)$ gives us $\pi_i(A) = \alpha_i\mu_i(A) + (1 - \alpha_i)\tilde{\overline{\pi}}_i(A)$, which implies that $\left(\tilde{\mathcal{E}},\, (\alpha_i)_{i\in I}\,,\, (\mu_i)_{i\in I}\right)$ is an echo chamber representation of $(\pi_i)_{i\in I}$.

A.2 Proof of Proposition 2

From the representation we have

$$\pi_i(A) = \alpha_i\mu_i(A) + (1 - \alpha_i)\overline{\pi}_{E(i)}(A)$$

Rearrange this to obtain

$$\frac{\pi_i(A)}{\overline{\pi}_{E(i)}(A)} = \alpha_i \frac{\mu_i(A)}{\overline{\pi}_{E(i)}(A)} + (1 - \alpha_i)$$

Since μ_i and $\overline{\pi}_{E(i)}$ are probabilities, $\frac{\mu_i(A)}{\overline{\pi}_{E(i)}(A)} \geq 0$, which implies $\alpha_i \geq 1 - \min \frac{\pi_i(A)}{\overline{\pi}_{E(i)}(A)}$. Since $\alpha_i \in (0,\, 1)$ by definition, we have the first part of the result.

For the second, choose α_i in the given set for each $i \in I$. Define μ_i as:

$$\mu_i(A) = \frac{\pi_i(A) - (1 - \alpha_i)\overline{\pi}_{E(i)}(A)}{\alpha_i}$$

Note that by the choice of α_i, we have that $\frac{\pi_i(A)}{\overline{\pi}_{E(i)}(A)} \geq 1 - \alpha_i$, which implies $\mu_i(A) \geq 0$ for A such that $\pi_i(A) \neq 0$. For $\pi_i(A) = 0$, we have that $\overline{\pi}_{E(i)}(A) = 0$ by proposition 1, which implies that $\mu_i(A) = 0$. By the argument in proposition 1, μ_i is additive, and the chosen parameters provide an echo chamber representation of $(\pi_i)_{i\in I}$.

A.3 Proof of Theorem 3

We first show that certainty conformism is sufficient for an echo chamber representation.

Define chambers as $E(i) = N(i) \cup \{i\}$. Since $j \in N(i)$ iff $i \in N(j)$, $j \in E(i)$ iff $E(i) = E(j)$. Then we have a valid partition of I. Certainty conformism implies that for all $i \in I$, $N(i) \neq \emptyset$. Thus, $|E(i)| \geq 2$ for all $i \in I$. By proposition 1, there exists an echo chamber representation under this partition.

To show necessity, note that in the absence of certainty conformism, there exists $j \in I$ such that $N(j) = \emptyset$. Then $|N(j) \cup \{j\}| = 1$. However, by proposition 1, if there exists an echo chamber representation with the partition $\mathcal{E}$, then $E(j) \subseteq N(j) \cup \{j\}$, which implies $|E(j)| \leq 1$. However, the definition of echo chambers requires that they are non-singleton, which is a contradiction.

A.4 Proof of Theorem 1

We first show that the axioms are sufficient for a maximal influence echo chamber representation.

Step 1. Defining chambers.

Define $R(i) = N(i) \cup \{i\}$ for each $i \in I$. By A1, $N(i) \neq \emptyset$, which means $|R(i)| \geq 2$ for all $i \in I$. Since $j \in R(i)$ iff $i \in R(j)$, setting $E(i) = R(i)$ gives us a valid partition of I, $\mathcal{E}$. By theorem 3, there exists an echo chamber representation under this partition.

Step 2. Defining core beliefs and immunity from influence parameters.

Given any event A, we can write

$$\begin{aligned}
\overline{\pi}_{R(i)}(A) &= \frac{\pi_i(A)}{|R(i)|} + \frac{|N(i)|}{|R(i)|}\overline{\pi}_{N(i)}(A) \\
\pi_i(A) &= |R(i)|\overline{\pi}_{R(i)}(A) - |N(i)|\overline{\pi}_{N(i)}(A) \\
\pi_i(A) &= (|N(i)| + 1)\overline{\pi}_{R(i)}(A) - |N(i)|\overline{\pi}_{N(i)}(A) \\
\pi_i(A) &= \overline{\pi}_{R(i)}(A) + |N(i)|(\overline{\pi}_{R(i)}(A) - \overline{\pi}_{N(i)}(A))
\end{aligned}$$

By A2, we know that there exists some event A such that $\overline{\pi}_{R(i)}(A) \neq \overline{\pi}_{N(i)}(A)$, and since $N(i) \neq \emptyset$, $\pi_i(A) \neq \overline{\pi}_{R(i)}(A)$. Then either $\pi_i(A) < \overline{\pi}_{R(i)}(A)$ or $\pi_i(A^c) < \overline{\pi}_{R(i)}(A)$. Consequently, $\min \frac{\pi_i(A)}{\overline{\pi}_{R(i)}(A)} < 1$. Choose $\alpha_i = 1 - \min \frac{\pi_i(A)}{\overline{\pi}_{R(i)}(A)} > 1 - 1 = 0$. Furthermore, $\pi_i(A) = 0$ iff $\pi_j(A) = 0$ for all $j \in R(i)$ iff $\overline{\pi}_{R(i)}(A) = 0$, which implies $\frac{\pi_i(A)}{\overline{\pi}_{R(i)}(A)} > 0$ for all A such that this ratio is defined. Then $\alpha_i < 1$, and $\alpha_i \in \left[1 - \min \frac{\pi_i(A)}{\overline{\pi}_{R(i)}(A)},\, 1\right) \cap (0,\, 1)$. By proposition 2, there exists an echo chamber representation with the partition $\mathcal{E}$ and immunity from influence parameters $(\alpha_i)_{i\in I}$. Rearrange the definition of behavioral beliefs in the model to get μ_i as

$$\mu_i(A) = \frac{\pi_i(A) - (1-\alpha_i)\overline{\pi}_{R(i)}(A)}{\alpha_i}$$

Then $\left(\mathcal{E},\, (\alpha_i)_{i\in I},\, (\mu_i)_{i\in I}\right)$ is an echo chamber representation of $(\pi_i)_{i\in I}$.

Step 3. Establishing that this representation is maximal influence.

By proposition 1, we know that if there exists an echo chamber representation of $(\pi_i)_{i\in I}$ under any partition $\tilde{\mathcal{E}}$, then $\tilde{E}(i) \subseteq R(i) = E(i)$. Thus there does not exist any representation with a coarsening of $\mathcal{E}$. If $\tilde{E}(i) = R(i)$ for all $i \in I$, note by proposition 2 that the corresponding immunity from influence parameter $\tilde{\alpha}_i \geq 1 - \min \frac{\pi_i(A)}{\overline{\pi}_{R(i)}(A)} = \alpha_i$. Then, to show that $\left(\mathcal{E},\, (\alpha_i)_{i\in I},\, (\mu_i)_{i\in I}\right)$, we must ensure that if there exists another echo chamber representation $\left(\tilde{\mathcal{E}},\, (\tilde{\alpha}_i)_{i\in I},\, (\tilde{\mu}_i)_{i\in I}\right)$, such that $\tilde{\mathcal{E}} = \mathcal{E}$ and $\tilde{\alpha}_i = \alpha_i$ for all $i \in I$, then $\tilde{\mu}_i = \mu_i$ for all $i \in I$.

By the definition of the model, for all $A \subseteq S$:

$$\begin{aligned}\tilde{\mu}_i(A) &= \frac{\pi_i(A) - (1-\tilde{\alpha}_i)\overline{\pi}_{R(i)}(A)}{\tilde{\alpha}_i} \\ &= \frac{\pi_i(A) - (1-\alpha_i)\overline{\pi}_{R(i)}(A)}{\alpha_i} \\ &= \mu_i(A)\end{aligned}$$

Thus $\mu_i = \tilde{\mu}_i$ for each $i \in I$. Then $\left(\tilde{\mathcal{E}},\, (\tilde{\alpha}_i)_{i\in I},\, (\tilde{\mu}_i)_{i\in I}\right)$ is identical to $\left(\mathcal{E},\, (\alpha_i)_{i\in I},\, (\mu_i)_{i\in I}\right)$, which means it cannot be non-maximal. Consequently, it is a maximal influence echo chamber representation of $(\pi_i)_{i\in I}$.

This completes the proof for sufficiency of the axioms.

We now look at the necessity of the two axioms.

Axiom A1: By theorem 3, certainty conformism is necessary for the existence of an echo chamber representation, of which the maximal influence echo chamber representations are a subclass.

Axiom A2: Suppose there exists $i \in I$ such that $\overline{\pi}_{R(i)}(A) = \overline{\pi}_{N(i)}(A)$ for all $A \subseteq S$.

Then note

$$\begin{aligned}\pi_i(A) &= \overline{\pi}_{R(i)}(A) + |N(i)| \left(\overline{\pi}_{R(i)}(A) - \overline{\pi}_{N(i)}(A)\right) \\ &= \overline{\pi}_{R(i)}(A)\end{aligned}$$

Then $\pi_i = \overline{\pi}_{R(i)}$. If A1 is violated, there does not exist any echo chamber representation of $(\pi_j)_{j\in I}$, so assume it is satisfied. Let $\left(\mathcal{E},\, (\alpha_i)_{j\in I}\,,\, (\mu_i)_{j\in I}\right)$ be a maximal influence echo chamber representation of $(\pi_j)_{j\in I}$. By proposition 1, we know that $E(i) = R(i)$. By definition, $\alpha_i \in (0,\, 1)$. Then there exists $\tilde{\alpha}_i \in (0,\, 1)$ such that $\tilde{\alpha}_i < \alpha_i$. Set $\mu_i = \pi_i$. Note that by $\tilde{\mu}_i = \pi_i = \overline{\pi}_{R(i)}$, the parameters $\tilde{\alpha}_i$ and $\tilde{\mu}_i$ imply $\pi_i(A) = \tilde{\alpha}_i\tilde{\mu}_i(A) + (1-\tilde{\alpha}_i)\overline{\pi}_{R(i)}(A)$ for all $A \subseteq S$.

Leaving the parameters associated with all others unchanged, note that $\pi_j(A) = \alpha_j\mu_j(A) + (1-\alpha_j)\overline{\pi}_{R(j)}(A)$ by supposition. Then $(\mathcal{E},\, (\tilde{\alpha}_i,\, \alpha_{-i})\,,\, (\tilde{\mu}_i,\, \mu_{-i}))$ is an echo chamber representation of $(\pi_j)_{j\in I}$. However, as $\tilde{\alpha}_i < \alpha_i$, $\left(\mathcal{E},\, (\alpha_i)_{j\in I}\,,\, (\mu_i)_{j\in I}\right)$ cannot be maximal influence, leading to a contradiction.

A.5 Proof of Theorem 4

Step 1. Show uniqueness of partition.

By theorem 1, if there exists a maximal influence echo chamber representation of $(\pi_i)_{i\in I}$, then $N(i) \neq \emptyset$, which implies that $|R(i)| = |N(i) \cup \{i\}\,| \geq 2$. Since $\mathcal{E}$, where $E(i) = R(i)$, is a partition of I, by proposition 1, there exists an echo chamber representation of $(\pi_i)_{i\in I}$ with the partition $\mathcal{E}$, and further, there does not exist $\tilde{\mathcal{E}}$, a coarsening of $\mathcal{E}$, such that there exists an echo chamber representation of $(\pi_i)_{i\in I}$ under $\tilde{\mathcal{E}}$.

Step 2. Show uniqueness of immunity from influence parameters.

Again, by the arguments in theorem 1, A2 implies that $\min \frac{\pi_i(A)}{\overline{\pi}_{R(i)}(A)} < 1$, which means that $1 - \min \frac{\pi_i(A)}{\overline{\pi}_{R(i)}(A)} > 0$ for all $i \in I$. Then, by proposition 2, there exists an echo chamber representation such that $\alpha_i = 1 - \min \frac{\pi_i(A)}{\overline{\pi}_{R(i)}(A)}$. Again, by proposition 2, if there exists some other representation with partition $\mathcal{E}$ and immunity from influence parameters $(\tilde{\alpha}_i)_{i\in I}$, $\tilde{\alpha}_i \geq 1 - \min \frac{\pi_i(A)}{\overline{\pi}_{R(i)}(A)} = \alpha_i$. Thus, for any maximal influence echo chamber representation, $\alpha_i = 1 - \min \frac{\pi_i(A)}{\overline{\pi}_{R(i)}(A)}$, which is unique.

Step 3. Show uniqueness of core beliefs.

In the proof of theorem 1, we have shown that if there exist two representations $\left(\mathcal{E},\, (\alpha_i)_{i\in I}\,,\, (\mu_i)_{i\in I}\right)$ and $\left(\tilde{\mathcal{E}},\, (\tilde{\alpha}_i)_{i\in I}\,,\, (\tilde{\mu}_i)_{i\in I}\right)$ such that $\tilde{\mathcal{E}} = \mathcal{E}$ and $\tilde{\alpha}_i = \alpha_i$ for each $i \in I$, then $\tilde{\mu}_i = \mu_i$ as well.

Since we have established that for any maximal influence echo chamber representation, the partition must be given by $\mathcal{E}$ with $E(i) = R(i)$ and $\alpha_i = 1 - \min \frac{\pi_i(A)}{\overline{\pi}_{R(i)}(A)}$ for each $i \in I$, it must be that μ_i is also uniquely determined. Thus, we have shown that the maximal influence echo chamber representation for any $(\pi_i)_{i \in I}$ is unique.

A.6 Proof of Theorem 2

First, we prove that an MIEC representation satisfies DFC-Inter and DFC-Intra. Let $\left(\mathcal{E}, (\alpha_i)_{i \in I}, (\mu_i)_{i \in I}\right)$ be an MIEC representation of $(\pi_i)_{i \in I}$. As established in the proof of theorem 4, under this representation $E(i) = N(i) \cup \{i\}$ and $\alpha_i = 1 - \min \frac{\pi_i(A)}{\overline{\pi}_{E(i)}(A)}$ for each $i \in I$. From remark 1

$$\overline{\pi}_{E(i)}(A) = \sum_{j \in E(i)} \frac{\alpha_j \mu_j(A)}{\sum_{k \in E(i)} \alpha_k}$$

Then for any $A \subseteq S$, $\pi_j(A) = 1$ for all $j \in E(i)$ iff $\overline{\pi}_{E(i)}(A) = 1$ iff $\mu_j(A) = 1$ for all $j \in E(i)$. Then take any two chambers E and E'. If $i \in E$ and $j \in E'$, note that $j \notin N(i)$, which means there exists A such that $\pi_i(A) = 1$ but $\pi_j(A) \neq 1$, or vice versa. Suppose w.l.o.g. that A^0 is such that $\pi_i(A^0) = 1$ but $\pi_j(A^0) < 1$. Then $\mu_{i'}(A^0) = 1$ for all $i' \in E$, which means E is fundamentally certain about A^0. However, $\pi_j(A^0) < 1$, which implies E' is not fundamentally certain about A^0. Since this is true for all E, E', DFC-Inter is satisfied.

Now by $\alpha_i = 1 - \min \frac{\pi_i(A)}{\overline{\pi}_{E(i)}(A)}$, suppose $A^0 \in \operatorname{argmin} \frac{\pi_i(A)}{\overline{\pi}_{E(i)}(A)}$. Then note that $\frac{\pi_i(A^0)}{\overline{\pi}_{E(i)}(A^0)} = 1 - \alpha_i = \alpha_i \frac{\mu_i(A^0)}{\overline{\pi}_{E(i)}(A^0)} + 1 - \alpha_i$. This implies $\mu_i(A^0) = 0$. However, $\overline{\pi}_{E(i)}(A^0) > 0$ as $\frac{\pi_i(A^0)}{\overline{\pi}_{E(i)}(A^0)}$ is defined, which implies that there exists $j \in E(i)$ such that $\mu_j(A^0) > 0$. Then i is fundamentally certain about $(A^0)^c$ but j is not. Since we can find such an A^0 for each $i \in I$, DFC-Intra is also be satisfied.

Now we show that if $\left(\mathcal{E}, (\alpha_i)_{i \in I}, (\mu_i)_{i \in I}\right)$ is a representation that satisfies DFC-Inter and DFC-Intra, it must be an MIEC representation.

By DFC-Inter, for any E, E', there exists an event A such that one chamber is fundamentally certain about it but the other is not. Suppose E is fundamentally certain about A and E' is not. By the earlier arguments, $\pi_i(A) = 1$ for all $i \in E$ and $\pi_j(A) < 1$ for all $j \in E'$. By proposition 1, $E(i) \subseteq N(i) \cup \{i\}$ for all i. On the other hand, if $i \in E$ and $j \in E'$, since $\pi_i(A) = 1$ and $\pi_j(A) < 1$, $j \notin N(i)$. Since this is true for all E, E', $N(i) \cup \{i\} \subseteq E(i)$. Thus, $E(i) = N(i) \cup \{i\}$, which means that there does not exist a coarsening of $\mathcal{E}$ under which partition there exists an EC representation of $(\pi_i)_{i \in I}$.

Then by DFC-Intra, there exists B for each $i \in I$ such that $\mu_i(B) = 1$ but $\mu_j(B) < 1$

for some $j \in E(i)$. Then $\mu_i(B^c) = 0$ and $\mu_j(B^c) > 0$, which means $\overline{\pi}_{E(i)}(B^c) > 0$. Then note that $\frac{\pi_i(B^c)}{\overline{\pi}_{E(i)}(B^c)} = 1 - \alpha_i$, which implies $\alpha_i = 1 - \frac{\pi_i(B^c)}{\overline{\pi}_{E(i)}(B^c)}$. Since $\frac{\pi_i(A)}{\overline{\pi}_{E(i)}(A)} \geq 1 - \alpha_i$ as $\mu_i(A) \geq 0$, $\frac{\pi_i(B^c)}{\overline{\pi}_{E(i)}(B^c)} = \min \frac{\pi_i(A)}{\overline{\pi}_{E(i)}(A)}$. By proposition 2, for any other representation $\left(\mathcal{E}, (\tilde{\alpha}_i)_{i \in I}, (\tilde{\mu}_i)_{i \in I}\right)$, $\tilde{\alpha}_i \geq \alpha_i$, which means that $\left(\mathcal{E}, (\alpha_i)_{i \in I}, (\mu_i)_{i \in I}\right)$ is an MIEC representation.

A.7 Proof of Proposition 3

Consider the problem of a single component. Since all individuals, $i \in E$, receive private information only up to some finite period, $\mu_E^t \equiv \mu_E$, for all $t \geq \max\{T_i | i \in E\}$, with $\mu_i^t = \mu_i^{T_i}$. Then

$$\pi_E^{t+1} = \Lambda_E \mu_E + |E|^{-1}(I_{|E|} - \Lambda_E)\mathbf{1}_{|E|}\pi_E^t$$

Define B as

$$B_E \equiv |E|^{-1}(I_{|E|} - \Lambda_E)\mathbf{1}_{|E|}$$

and write

$$\begin{aligned} \pi_E^{t+1} &= \Lambda_E \mu_E + B_E \pi_E^t \\ &= (I_{|E|} + B_E + B_E^2 + \dots + B_E^t)\Lambda_E \mu_E + B_E^{t+1}\pi_E^0 \end{aligned}$$

It can easily be verified that $||B_E|| \leq \max_{i \in E}(1 - \alpha_i)$, which implies that

$$\lim_{t \to \infty} \pi_E^t = \left(\sum_{k=0}^{\infty} B_E^k\right) \Lambda_E \mu_E + \lim_{t \to \infty} B_E^{t+1}\pi_E^0$$

Since $||B_E|| < 1$, $\lim_{t \to \infty} B_E^{t+1} = 0$, and

$$\begin{aligned} \lim_{t \to \infty} \pi_E^t &= (I_{|E|} - B_E)^{-1}\Lambda_E \mu_E \\ &= (I_{|E|} - |E|^{-1}(I_{|E|} - \Lambda_E)\mathbf{1}_{|E|})^{-1}\Lambda_E \mu_E \end{aligned}$$

irrespective of the initial condition on reported beliefs. Now we must simply show that $\lim_{t\to\infty} \pi_E^t$ is a fixed point of this process. Let $\pi_E^* = \lim_{t\to\infty} \pi_E^t$. Note that

$$\begin{aligned}
\Lambda_E \mu_E + B_E \pi_E^* &= \Lambda_E \mu_E + B_E (I_{|E|} - B_E)^{-1} \Lambda_E \mu_E \\
&= (I_{|E|} + B_E (I_E - B_E)^{-1}) \Lambda_E \mu_E \\
&= \left(I_E + B_E \sum_{k=0}^{\infty} B_E^k \right) \Lambda_E \mu_E \\
&= \left(\sum_{k=0}^{\infty} B_E^k \right) \Lambda_E \mu_E \\
&= (I_{|E|} - B_E)^{-1} \Lambda_E \mu_E \\
&= \pi_E^*
\end{aligned}$$

Then, for each $i \in E$ and $A \subseteq S$,

$$\pi_i^*(A) = \lambda_i \mu_i(A) + (1 - \lambda_i) \frac{1}{|E|} \sum_{j \in E} \pi_j^*(A)$$

Since this is true for any arbitrary component, the argument goes through for all components in $\mathcal{E}$. However, Remark 1 established that $(\pi_i)_{i\in I}$, the behavioral beliefs corresponding to the parameters $\mathcal{E}$, $(\mu_i)_{i\in I}$, and $(\alpha_i)_{i\in I}$ with $\alpha_i = \lambda_i$ is the unique solution to the above equation, implying that $\pi_i^* = \pi_i$ for all $i \in I$.

A.8 Proof of Proposition 4

Proposition 4 can be proved as follows.

$$\begin{aligned}
\pi_i(\sigma_i(c^i) \,|\, c^i) = 1 &\iff \overline{\pi}_i(\sigma_i(c^i) \,|\, (c_1, ..., c_k)) = 1 \\
&\iff \mu_j(\sigma_i(c^i) \,|\, c^j) = 1 \; \forall \, j \in E(i)
\end{aligned}$$

Take some state s such that $s \in \sigma_i(c^i)$. Then $\mu_i(s \,|\, c^i) > 0$. This means, however, that $\pi_j(s \,|\, c^j) > 0$ for all $j \in E(i)$. Then, if $\pi_j(\sigma_j(c^j) \,|\, c^j) = 1$ for all $j \in E(i)$, it must be that $s \in \sigma_j(c^j)$. This implies that $\sigma_i(c^i) \subseteq \sigma_j(c^j)$ for all $j \in E(i)$. By a symmetric argument, $\sigma_j(c^j) \subseteq \sigma_i(c^i)$, and we can extend this to any $i, j \in E_k$ some echo chamber.

Let us now assume that $\sigma_j(c^j) = \sigma_i(c^i)$ for all $i, j \in E_k$. Then $\mu_i(\sigma_i(c^i)|c^i) = \mu_i(\sigma^j(c^j)|c^i) = 1$ for all $i, j \in E_k$. Then for all $i \in E_k$, $\mu_j(\sigma_i(c^i)|c^j) = 1$ for all $j \in E_k$. Thus, $\pi_i'(\sigma_i(c^i)) = 1$ for all $i \in E_k$. This completes the proof.

A.9 Proof of Lemma 1

Note that if $i \in L$, then upon receiving the signal c_r, the updated behavioral belief of i is given by

$$\begin{aligned}\pi_i'(\ell) &= \alpha\mu_L(\ell|c_r) + (1-\alpha)\left[\frac{X_\ell^L}{n}\mu_L(\ell|c_\ell) + \frac{n-X_\ell^L}{n}\mu_L(\ell|c_r)\right] \\ &= \frac{(1-\rho)q}{(1-\rho)q+\rho(1-q)} + (1-\alpha)\frac{X_\ell^L}{n}\left[\frac{\rho q}{\rho q+(1-\rho)(1-q)} - \frac{(1-\rho)q}{(1-\rho)q+\rho(1-q)}\right]\end{aligned}$$

where X_ℓ^L is the random variable that specifies the number of individuals in E_L who have received the signal c_ℓ. Then $\pi_i'(\ell) < \frac{1}{2}$ iff

$$\begin{aligned}&\frac{1}{2} > \frac{(1-\rho)q}{(1-\rho)q+\rho(1-q)} + (1-\alpha)\frac{X_\ell^L}{n}\left[\frac{\rho q}{\rho q+(1-\rho)(1-q)} - \frac{(1-\rho)q}{(1-\rho)q+\rho(1-q)}\right] \\ \iff\ &\frac{X_\ell^L}{n} < \frac{1}{1-\alpha}\underbrace{\frac{\frac{1}{2}-\frac{(1-\rho)q}{(1-\rho)q+\rho(1-q)}}{\frac{\rho q}{\rho q+(1-\rho)(1-q)}-\frac{(1-\rho)q}{(1-\rho)q+\rho(1-q)}}}_{\equiv x^*}\end{aligned}$$

$\alpha \leq \frac{1}{2}$ implies $(1-\alpha)^{-1} \leq 2$. Note that x^* is increasing in ρ. If $\rho > q$, by Remark 3, x^* is positive and bounded by half, which implies $(1-\alpha)^{-1}x^* \leq 1$. Since $\rho \neq 1$, $x^* < \frac{1}{2}$, so $\hat{x}(\alpha, \rho, q) = (1-\alpha)^{-1}x^* < 1$.

If $\rho \leq q$, note that $\mu_L(\ell|c) \geq \frac{1}{2}$ irrespective of whether $c = c_\ell$ or $c = c_r$. Then for all realizations of signals, $\pi_i'(\ell) \geq \frac{1}{2}$. Let $\hat{x}(\alpha, \rho, q) = 0$. Note that $\frac{X_\ell^L}{n} \geq 0$ for all realizations of signals. Then $\pi_i'(\ell) < \frac{1}{2} \iff \frac{X_\ell^L}{n} \leq \hat{x}(\alpha, \rho, q)$ vacuously.

A.10 Proof of Lemma 2

Note from the proof of Lemma 1 that $x^* \in [0, \frac{1}{2}]$ and x^* is continuous in ρ. Since $x^* = 0$ when $\rho = q$ and $x^* = \frac{1}{2}$ when $\rho = 1$, and $0 < 1-\alpha < \frac{1}{2}$, by intermediate value theorem, there exists $\rho^* \in (q, 1)$ such that $x^* = 1-\alpha$ when $\rho = \rho^*$.

A.11 Proof of Proposition 5

Since we assume that $q < \rho < \rho^*(\alpha, q)$, $\hat{x}(\alpha, \rho, q) < 1$. Assume n large enough such that $\frac{n-1}{n} > \hat{x}(\alpha, \rho, q)$.

Since the problem is symmetric, suppose the true state is ℓ. We want to analyze all situations where voting based on core and behavioral beliefs result in different policies being implemented. Policy R is implemented under behavioral beliefs if either (i) $\frac{X_r^R}{n} \geq \hat{x}(\alpha, \rho, q)$ and $\frac{X_\ell^L}{n} < \hat{x}(\alpha, \rho, q)$ or (ii) $\frac{X_\ell^L}{n} < \frac{X_r^R}{n} < \hat{x}(\alpha, \rho, q)$. A symmetric construction works for when policy L is implemented. In both cases, note that if a policy is implemented under behavioral beliefs, then it must be implemented under core beliefs as well. The only remaining case is when a certain policy may be implemented when voting by core beliefs, but voting under behavioral beliefs is deadlocked, meaning either policy may be voted in with probability $\frac{1}{2}$.

We only need to establish, then, that for all situations such that voting under behavioral beliefs is deadlocked, it is more likely that voting based on core beliefs elects the optimal policy. To see this, note that policy R is implemented under core beliefs iff $X_r^R > X_\ell^L$. Voting under behavioral beliefs is deadlocked under such a scenario iff $X_r^R > X_\ell^L \geq n\hat{x}(\alpha, \rho, q)$. Similarly, for policy L, with $X_\ell^L > X_r^R$.

Since $\rho \geq \frac{1}{2}$, if $a > b$, $\rho^{a+(n-b)}(1-\rho)^{b+(n-a)} > \rho^{b+(n-a)}(1-\rho)^{a+(n-b)}$. If $X_\ell^L = a$ and $X_r^R = b$ and $a > b \geq n\hat{x}(\alpha, \rho, q)$, then policy L is implemented under voting by core beliefs with a deadlock arising from voting under behavioral beliefs. However, for $X_r^R = a$ and $X_\ell^L = b$, policy R is implemented under voting by core beliefs, without changing the outcomes under behavioral beliefs. However, note that $\mathbb{P}(X_r^R = a,\, X_\ell^L = b) < \mathbb{P}(X_\ell^L = a,\, X_r^R = b)$, which means that the probability of voting for policy R by core beliefs is less than half. On the other hand, policy R is voted in half of the time under behavioral beliefs. Thus, we have shown $\gamma(\alpha, \rho, q) > \tilde{\gamma}(\alpha, \rho, q)$.

Now, when $\hat{x}(\alpha, \rho, q) \to 1$, note that for any n, $\frac{n-1}{n} < \hat{x}(\alpha, \rho, q)$ after a threshold. Thus, voting outcomes under core and behavioral beliefs are identical. If $\hat{x}(\alpha, \rho, q) \to 0$, then $\rho \to q$, in which case voting outcomes converge to being deadlocked almost surely, whether under core or behavioral beliefs.

A.12 Proof of Proposition 6

By the conclusions of Lemmas 1 and 2, there exists $\hat{x}(\alpha, \rho, q) \in [0, 1]$ such that $\pi_i'(\ell) < \frac{1}{2}$ for $i \in L$ conditional on receiving the signal c_r iff $\frac{X_\ell^L}{n} < \hat{x}(\alpha, \rho, q)$. We also noted that for $\rho \in (q, \rho^*(\alpha, q))$, $\hat{x}(\alpha, \rho, q)$ is continuous and strictly increasing in ρ.

Note, now, that by the strong law of large numbers, $\frac{X_\ell^L}{n} \to 1-\rho$ almost surely as $n \to \infty$. This implies that $\pi_i'(\ell) < \frac{1}{2}$ almost surely as $n \to \infty$ iff $1-\rho < \hat{x}(\alpha, \rho, q)$. Suppose $q < 1$. Since $\inf_\alpha \hat{x}(\alpha, 1, q) = \frac{1}{2} > 1 - \rho \equiv 0$, and $\hat{x}(\alpha, q, q) = 0 < 1 - q \equiv 1 - \rho$, there exists a root of $\hat{x}(\alpha, \rho, q) + \rho - 1$ for $\rho \in (q, 1)$. That value of ρ becomes $\tilde{\rho}(\alpha, q)$. On the other hand, if $q = 1$, then $\pi_i'(\ell)$ is never less than $\frac{1}{2}$, which means that $\tilde{\rho}(\alpha, q) = 1$.

References

Acemoglu, Daron, Kostas Bimpikis, and Asuman Ozdaglar. 2014. "Dynamics of information exchange in endogenous social networks." *Theoretical Economics* 9 (1): 41–97.

Acemoglu, Daron, Munther A. Dahleh, Ivan Lobel, and Asuman Ozdaglar. 2011. "Bayesian Learning in Social Networks." *The Review of Economic Studies* 78 (4): 1201–1236.

Azomahou, Theophile T., and Daniel C. Opolot. 2014. *Beliefs dynamics in communication networks.* Working Paper, UNU-MERIT Working Paper Series 2014-034. Maastricht: Maastricht Economic, Social Research Institute on Innovation, and Technology.

Baccara, Mariagiovanna, and Leeat Yariv. 2016. "Choosing peers: Homophily and polarization in groups." *Journal of Economic Theory* 165 (September): 152–178.

Bala, Venkatesh, and Sanjeev Goyal. 1998. "Learning from Neighbours." *Review of Economic Studies* 65 (3): 595–621.

Barberá, Pablo. 2020. "Social Media, Echo Chambers, and Political Polarization." In *Social Media and Democracy: The State of the Field and Prospects for Reform,* edited by Joshua A Tucker and Nathaniel Persily, 34–55. Cambridge University Press.

Borah, Abhinash, and Christopher Kops. 2018. "Choice via social influence." *Unpublished manuscript.*

Chambers, Christopher P, Tugce Cuhadaroglu, and Yusufcan Masatlioglu. 2019. "Behavioral influence." *Unpublished manuscript.*

Chambers, Christopher P, Yusufcan Masatlioglu, and Christopher Turansick. 2021. "Influence and Correlated Choice." *arXiv preprint arXiv:2103.05084.*

Condorcet, Marquis de. 1785. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix.* Paris: L'Imprimerie Royale.

Cookson, J Anthony, Joseph E Engelberg, and William Mullins. 2023. "Echo Chambers." *The Review of Financial Studies* 36 (2): 450–500.

Cuhadaroglu, Tugce. 2017. "Choosing on influence." *Theoretical Economics* 12 (2): 477–492.

DeGroot, Morris H. 1974. “Reaching a Consensus.” *Journal of the American Statistical Association* 69 (345): 118–121.

DeMarzo, P. M., D. Vayanos, and J. Zwiebel. 2003. “Persuasion Bias, Social Influence, and Unidimensional Opinions.” *The Quarterly Journal of Economics* 118 (3): 909–968.

Dietrich, Franz. 2008. “The Premises of Condorcet’s Jury Theorem Are Not Simultaneously Justified.” *Episteme* 5 (1): 56–73.

Eyster, Erik, and Matthew Rabin. 2010. “Naïve Herding in Rich-Information Settings.” *American Economic Journal: Microeconomics* 2 (4): 221–243.

Fershtman, Chaim, and Uzi Segal. 2018. “Preferences and Social Influence.” *American Economic Journal: Microeconomics* 10 (3): 124–142.

Friedkin, Noah E., and Eugene C. Johnsen. 1990. “Social influence and opinions.” *The Journal of Mathematical Sociology* 15 (3-4): 193–206. ISSN: 0022-250X. https://doi.org/10.1080/0022250X.1990.9990069.

Glaeser, Edward L., and Cass R. Sunstein. 2009. “Extremism and Social Learning.” *Journal of Legal Analysis* 1 (1): 263–324.

Golub, Benjamin, and Matthew O Jackson. 2010. “Naïve Learning in Social Networks and the Wisdom of Crowds.” *American Economic Journal: Microeconomics* 2 (1): 112–149.

Jiménez-Martínez, Antonio. 2015. “A model of belief influence in large social networks.” *Economic Theory* 59 (1): 21–59.

Kaniovski, Serguei, and Alexander Zaigraev. 2011. “Optimal jury design for homogeneous juries with correlated votes.” *Theory and Decision* 71 (4): 439–459.

Kashaev, Nail, and Natalia Lazzati. 2019. “Peer effects in random consideration sets.” *arXiv preprint arXiv:1904.06742.*

Lazzati, Natalia. 2020. “Codiffusion of Technologies in Social Networks.” *American Economic Journal: Microeconomics* 12 (4): 193–228.

Levy, Gilat, and Ronny Razin. 2015a. “Correlation Neglect, Voting Behavior, and Information Aggregation.” *American Economic Review* 105 (4): 1634–1645.

———. 2015b. “Does Polarisation of Opinions Lead to Polarisation of Platforms? The Case of Correlation Neglect.” *Quarterly Journal of Political Science* 10 (3): 321–355.

———. 2018. “Information diffusion in networks with the Bayesian Peer Influence heuristic.” *Games and Economic Behavior* 109 (May): 262–270.

Levy, Gilat, and Ronny Razin. 2019. "Echo Chambers and Their Effects on Economic and Political Outcomes." *Annual Review of Economics* 11 (1): 303–328.

Martinez, Gabriel, and Nicholas H. Tenev. 2024. *Optimal Echo Chambers,* October.

McCarty, N, K T Poole, and H Rosenthal. 2006. *Polarized America: The Dance of Ideology and Unequal Riches.* MIT Press.

Peleg, Bezalel, and Shmuel Zamir. 2012. "Extending the Condorcet Jury Theorem to a general dependent jury." *Social Choice and Welfare* 39 (1): 91–125.

Reshidi, Përllumb. 2024. "Information Sequencing with Naive Social Learning." *Unpublished Manuscript* (June). Accessed December 25, 2024. https://www.preshidi.com/papers/Information_Sequencing_with_Naive_Social_Learning.pdf.