Teorema, 2025, 44(1), pp. 1-20
https://doi.org/10.30827/trif.34113

How to Build an Epistemic Echo Chamber

Steven Hales

Commonwealth University at Bloomsburg

Abstract: The present paper offers an explanatory model of the logical structure of epistemic echo chambers. Most work done so far on echo chambers attributes their cause to epistemic vice: they arise due to poor reasoning or cognitive biases. I argue that instead they are a predictable outcome of rational belief formation. They result from groups aggregating and magnifying the beliefs of their members, and the reflection of group consensus back to individuals who then increase their confidence to better conform with the group. This process is recursive, driving both individuals and epistemic groups to complete conviction. I also argue that not all echo chambers are bad. They are virtuous when they increase credence in the truth for those inside and inoculate them against the falsehoods, corrupt data, and fallacious reasoning outside the chamber. Echo chambers are bad when they start with falsehoods and amplify confidence in them.

Keywords: *echo chambers, epistemic groups, suboptimal equilibria, Condorcet's jury theorem.*

Resumen: El presente artículo ofrece un modelo explicativo de la estructura lógica de las cámaras de eco epistémicas. La mayoría de los trabajos realizados hasta la fecha sobre cámaras de eco atribuyen su causa a vicios epistémicos: surgen debido a un razonamiento deficiente o sesgos cognitivos. Argumento que, en cambio, son un resultado predecible de la formación racional de creencias. Resultan de la agregación y amplificación de las creencias de sus miembros por parte de los grupos, y de la repercusión del consenso grupal en los individuos, quienes aumentan su confianza para integrarse mejor en el grupo. Este proceso es recursivo, impulsando tanto a los individuos como a los grupos epistémicos a una convicción completa. También argumento que no todas las cámaras de eco son malas. Son virtuosas cuando aumentan la credibilidad en la verdad de quienes las componen y las protegen contra las falsedades, los datos corruptos y el razonamiento falaz externo. Las cámaras de eco son malas cuando parten de falsedades y amplifican la confianza en ellas.

Palabras clave: *cámaras de eco, grupos epistémicos, equilibrios subóptimos, teorema del jurado de Condorcet.*

Received: 4 June 2025 **Accepted**: 22 June 2025 **Published**: 1 July 2025

Teorema. Revista Internacional de Filosofía
ISSN/ISSN-e: 1888-1254

* This work has been supported by Research Project PID2022-142120NB-100, financed by the Ministry of Economy and Competitiveness, Spanish Government.

1. Two ways to explain failure

It is easy to assume that if something went wrong, then someone did something wrong; false beliefs and poor outcomes must be the result of our fallible and error-prone nature. As Kant famously wrote, out of the crooked timber of humanity nothing entirely straight can ever be made. Bad results mean that mistakes were made, typically due to fallacious reasoning or cognitive biases. Fallacious reasoning is difficult to spot because it tends to be a small variation of perfectly legitimate reasoning; denying the antecedent is awfully similar to modus tollens, and having accepted the validity of the latter one might uncritically assume the former is valid too. Similarly cognitive biases are resilient because they are psychologically compelling. For example, it's hard to actively seek out counterevidence for propositions we want to be true, so we are prone to the confirmation bias. Likewise vivid and memorable stories are more psychologically salient than cold statistics, so we succumb to the availability error.

Fallacies and biases seem like the kinds of things that can be overcome by diligent mental hygiene. We know those inference traps can spring on the unwary, so we learn logic, statistics, and caution. However, there are also examples in which perfectly rational action still leads to suboptimal results, and these cases are more troubling because it is much less clear what to do about them.

For instance, there can be exogenous forces and circumstances that determine unfortunate outcomes, despite the best efforts of cooperating agents. One example is traffic jams. Some traffic jams are the result of easily identifiable causes, like an accident, road damage, or a fallen tree. Others, known as phantom traffic jams, are more mysterious—traffic crawls to a near halt, then speeds up, then slows back down again for no discernible reason. Recent work suggests that not only do traffic phenomena (density, jams) propagate in waves, but they are similar to explosion detonation waves and can be modeled by the same equations. Small perturbations in traffic flow amplify into phantom jams despite no "individual drivers behaving in a 'wrong' way. In fact [jams] can even occur if all drivers behave by the exact same laws" (Flynn et al., 2023). Everyone is driving the best they can and still phantom traffic jams emerge out of nowhere.

Other less than ideal consequences can arise from rational action when certain kinds of incentive structures are in place. The Prisoners' Dilemma is the most famous, and it is intractable precisely because it is not a failure of rationality. Defect/defect is the sole Nash equilibrium, it strictly dominates all other strategies, and, even when the players fully understand that cooperate/cooperate is Pareto-optimal, is the most reasonable choice for players in the game.

Another example is Malthusian traps, or races to the bottom. In these cases agents sacrifice a common value in order to increase their competitiveness. At a football game everyone is trying to see the game as best they can, peering over and around those in the seats closer to the field. I beat out others at this competition when I stand up to watch, although that means giving up my comfortable seat. By doing so I make it much harder for the person behind me to see the game, so they also stand up to compensate. In the end everyone behind the front row is standing, all having sacrificed the common value of being able to sit comfortably, and no one is able to see the game better than if everyone had remained seated.

The same trap is at work with two-earner families. Suppose all families are

one-earner households. My spouse already works, but I realize that we can get ahead financially if I get a job too. Both of us having paid work outside the home makes it much harder to mind children, take care of the housework, run errands, go shopping, meet contractors, and the other essential tasks of adult life. I make that sacrifice and take a job. Our family does prosper financially, at least until everyone else starts making the same choice. Then the costs of housing, transportation, and utilities go up as the average household supply of money goes up. In the end every family requires two earners to afford the higher cost of living while remaining in the same fiscal position relative to others, and everyone has given up the shared value of easily managing home life.

When we get unacceptable results (false beliefs, phantom traffic jams, Malthusian traps, suboptimal equilibria[1]), there are two broad-brush explanatory strategies for how we got here: (1) we screwed up, and (2) it couldn't be helped. We screwed up when those results are due to fallacious reasoning, cognitive biases, or other mistakes. Let's call this sort of explanation *Vice Theory*. Bad results couldn't be helped when they are the inexorable consequences of game-theoretic or decision-theoretic logic. Let's call this kind of explanation *Reasonable Expectations*. Vice Theory is the usual way that philosophers have addressed troubling phenomena in social epistemology such as echo chambers, epistemic bubbles, and belief polarization. In this paper I will build a Reasonable Expectations explanatory model that shows how echo chambers are a predictable, normal outcome of rational belief formation. The present paper is not a defense of echo chambers; it is a modeling of their construction.

The tasks ahead are first, describing how echo chambers are an expected outcome of rational belief formation; second, demonstrating that echo chambers can be virtuous when they increase credence in the truth for those inside and inoculate them against the falsehoods, corrupt data, and poor reasoning outside the chamber; third, showing how echo chambers go wrong when they magnify falsehoods within, and thwart those inside from attaining the truth; fourth, arguing that modeling echo chambers in this way has explanatory benefits that the vice theory approach does not.

2. Modeling echo chambers

What is an echo chamber? Kathleen Jamieson and Joseph Cappella write, "[By an 'echo chamber'] we mean to suggest a bounded, enclosed media space that has the potential to both magnify the messages delivered within it and insulate them from rebuttal" (Jamieson and Cappella, 2010) p. 76. Bert Baumgaertner avers, "An echo chamber is a sociological setting where peoples' prior beliefs are 'echoed back' giving the impression that their beliefs are correct" (Baumgaertner, 2014). According to Kenneth Boyd, "Echo chambers reinforce a set of views… in an echo chamber individuals tend to increase their confidence in their beliefs" (Boyd, 2019). C. Thi Nguyen elaborates, "I use 'echo chamber' to mean an epistemic community which creates a significant disparity in trust between members and non-members. This disparity is created by excluding non-members through epistemic discrediting while simultaneously amplifying members' epistemic credentials" (Nguyen, 2020). Endre Begby describes an echo chamber as "a structurally reinforced socio-epistemic mechanism that evinces some degree of active resistance to contrary opinions" (Begby 2022). Christopher Ranalli and Finlay Malcolm

[1] (Cf. Yudkowsky, 2017) for several other examples.

write, "Echo chambers distribute and reiterate information, where the content of that information is often consonant with the relevant in-group beliefs and other shared in-group attitudes of the people receiving it" (Ranalli and Malcolm, 2023).

The approach in what follows is to take the preceding informal, intuitive statements about echo chambers and develop a more precise model of how such echo chambers form. As with all such models, there is some idealization involved. For example, Boyle's Gas Law, which states that the pressure and volume of a gas are inversely proportional, treats molecules as frictionless spheres and ignores their electrochemical interactions. While a perfect empirical account of gas behavior would not make those assumptions, Boyle's Law was a great leap forward in understanding how gasses work. In fact, such idealizations are integral to understanding complex phenomena, and future empirical work is expected to be a refinement of the basic model (cf. Elgin 2017). Once my own positive account of echo chambers is developed, I will consider the chief competitor, namely Vice Theory, and argue that it is inferior to my approach for two reasons: (1) Vice Theory mistakenly treats all echo chambers as vicious and everyone within an echo chamber as unfairly prejudiced against external views, and (2) Vice Theory does not predict echo chambers becoming more prominent or polarization becoming more extreme whereas my account does. Therefore the model to be developed here is a more fruitful approach.

The most salient feature of echo chambers is that they are a mutual interplay between individual and group beliefs, where the conviction level of one is mirrored and augmented by the other (cf. Lackey, 2021a p. 207). While Nguyen, Jamieson, and Cappella are right that those inside an echo chamber dismiss and distrust those outside, that is not a constitutive feature of echo chambers. I will argue that it is a natural consequence of them.

In the Reasonable Expectations model there are three essential properties of echo chambers. Echo chambers are

1. *Amplifying*. Individual beliefs are aggregated and magnified by the group. This is the chamber.

2. *Reflecting*. Group consensus is reflected back to individual members, increasing their confidence. This is the echo.

3. *Recursive.* Once individuals raise their confidence to better conform with the group, those fresh credences are again amplified. This is how beliefs are repeated, reinforced, and steadily strengthen.

Obviously a full discussion of epistemic groups is beyond this paper. I am assuming, with Jennifer Lackey and others, that it is sensible to talk about group attitudes, including credences and belief (cf. Lackey, 2021b). I am also assuming this modest requirement: S is a member of an epistemic group only if S believes there is a >50% chance that the group is correct about >50% of the beliefs it holds within its domain. For example, Kris is a member of the epistemic group of biologists only if she thinks it is better than a coin toss that the group is right about most biological beliefs it holds. We cannot just say she thinks, of every biological belief held by the group, that there is a >50% chance the belief is right. Even a professional biologist could have a couple of outlier, heretical beliefs, or think the group got it wrong about one thing or another. But if most of her biological beliefs are heretical ones, Kris is not a member of

that epistemic group.

Similarly, someone is a Lutheran only if they think Lutherans are probably right about most religious claims. A Lutheran might disagree about this or that part of the doctrine, but if they think Lutherans are probably wrong about most religious issues, they are not a Lutheran. Again, I am only specifying a *necessary* condition for group membership. Other necessary conditions for being Lutheran could include baptism or catechism in the church. Clearly those are not sufficient conditions, though, as someone both baptized and catechized could subsequently become an atheist and reject all religion. General conformity with group belief is also required. Another way to see this is to compare Lutheranism and Episcopalianism. Both are mainstream Christian denominations and share a great number of religious beliefs. For someone to be determinately Lutheran but not determinately Episcopalian, she would have to accept most of the Lutheran beliefs when it comes to doctrinal differences with the Episcopalians.

As a last point of clarification, someone might culturally identify with a group without being a member of the associated epistemic group, as with someone who claims Jewish heritage but rejects Jewish beliefs. The focus here is only on epistemic groups.

3. How to get amplification

Echo chambers are amplifying in that they take individual beliefs and magnify their strength at the group level. For instance, as individuals, most libertarians may be positively disposed to small government, but as a group they are strongly supportive of it. To move from a set of individual attitudes to group-level attitudes requires an aggregation function. There are many possible aggregation functions, such as the group believing that P just in case a supermajority of members does, or all of the members do, or the Supreme Leader does. In fact, there is an extraordinarily large number of options: Christian List and Philip Pettit (List and Pettit, 2011 p. 50) calculate that if an attitude is to be formed on just one proposition by a ten member group, there are 2^{1024} logically possible aggregation functions. Many of these are obviously terrible, or lead to inconsistencies. However, since we cannot count on echo chambers to produce consistent attitudes, it will not matter whether the operant aggregation function admits of contradictory results.

To build an echo chamber we will need a means of aggregating individual credences that amplifies them at the group level.[2] There are many aggregation functions that will achieve this aim; in fact, any function that takes a group majority (simple, supermajority, unanimity, etc.), where most of the group gives a greater than .5 credence to P, will yield a stronger group credence on P. Here is how we can begin to build an echo chamber with simple majority voting in conjunction with Condorcet's Jury Theorem. Condorcet's theorem states: assume (Independence) that individual voters have independent probabilities of voting for the correct alternative. Also assume (Competence) that these probabilities exceed ½ for each voter. It follows that as the

[2] The argument of this paper is presented mostly in credence language, but that is not a requirement of the account. A good review of the relationship between belief and credence is (Jackson, 2020).

size of the group of voters increases, the probability of a correct majority increases and tends to one (infallibility) in the limit (Dietrich and Spiekermann, 2020). Condorcet originally used this result to show that a jury was much likelier to get the correct verdict than the average member, so long as Independence and Competence were satisfied.

Condorcet's jury theorem is the basis of the wisdom of crowds literature and discussed in democratic voting theory with respect to whether we should expect democracies to produce good decisions.[3] The Independence and Competence conditions are needed for Condorcetian voters to tend to produce the truth. A good bit of the technical literature surrounding the jury theorem focuses on how plausible Independence and Competence are under real-world conditions, to what extent they can be weakened, whether they are in tension with each other, the difference between causal and statistical independence, and so on (Dietrich and Spiekermann, 2020). Independence and Competence do not matter for the analysis of echo chambers, because the phenomenon under discussion is not whether echo chambers tend towards the truth, but how individuals and groups interact to produce agreement that is amplifying. To achieve amplification, all that is needed from Condorcet is the mathematics behind majority voting. In what follows, I will refer to "Condorcet's Jury Theorem," but really I will be relying on something more like "Condorcet Lite" that does not assume Competence or Independence.

Condorcet's Jury Theorem is surprising, so here is an intuitive way to view it.[4] Suppose you have an unfair coin. In fact, when you flip it, it comes up heads 60% of the time. Suppose you and your sister decide to flip the coin for the last piece of cake. Being smart, you call heads. There is still a 40% chance your sister gets the cake. Now suppose that instead of a single-flip game, you decide to award the cake on two out of three flips. Heads will win if flip 1=heads and flip 2=heads, or flip 1=heads and flip 3=heads, or flip 2=heads and flip 3= heads. The probability that heads will win in this game is 64.8%. Proof: pr(HH) + pr(HTH) + pr(THH) = (.6*.6) + (.6*.4*.6) + (.4*.6*.6)= .648. Conversely, the probability that tails will win is pr(TT) + pr(THT) + pr(HTT) = (.4*.4) + (.4*.6*.4) + (.6*.4*.4) = .352. Of course, .648 + .352= 1. There is a better chance of winning with heads in the two out of three game than with a single toss. What about playing three out of five? In this case there are 10 ways for heads to win, with a total probability of 68.256%. In single toss, the chance of heads winning was .6; in two out of three it was .648, in three out of five it was .68256. The larger the number of tosses the likelier the chance that heads wins. In the infinite limit, the probability that heads wins is 1. Analogously, the more jurors on Condorcet's jury that tend to vote in one direction, the more likely it is that the jury as a whole will vote that way.

The real power of Condorcet is that it drives to unanimity in group judgment. If most members of a group lean towards P, the group will conclude that P with a higher probability than the average member. The larger the group (as with the larger number of coin tosses), the more probable it will judge P to be. For example, if most Democrats think it is more likely than not that gun control is a good idea, your neighborhood Democrats will collectively think it is a very good idea, and Democrats as a whole will think it nearly certain that gun control is a good idea.

---

3 (Cf. Sunstein, 2006) and critical discussion in (Estlund, 2008 Ch. 12).

4 (Cf. Pettit, 2006 p. 158), and (Estlund, 2008 p. 224).

The first property of echo chambers is that they are *amplifying*: individual credences are aggregated such that group-level credence is higher. Condorcet's jury theorem shows how to get amplification: it produces group confidence that exceeds average individual confidence. Now we have groupthink, but not yet echo chambers. The second property of echo chambers is that they are *reflecting*. For that we need a means of mirroring group consensus back to believers, increasing their confidence. First I will provide psychological evidence that individuals do in fact increase their credence as a result of group beliefs, and then I will provide philosophical reasons to think that is a reasonable thing to do.

4. How to get reflection

Psychologists have long made the descriptive case that individuals routinely alter their judgments to conform with those of a group. In Solomon Asch's classic experiments in the 1950s, about a third of experimental subjects judged lines of obviously different lengths to be the same length when a majority group did so. In Asch's words, "many subjects apparently shifted their judgments in the direction of the views of the majorities" (Asch, 1955). Particularly striking is that the individual subjects were not in any important way part of the majority group; the others were merely in the same room. Even so, many subjects modified their perceptual beliefs to more closely hew to those of others. There have been numerous variants on and replications of Asch's studies, and the basic findings have held up, *pace* some gender and cultural variation (Bond and Smith, 1996). There is also recent evidence that conformity is strong when the group is rewarded for getting the correct answer, but conformity is weak when individuals are rewarded for getting it right (Fujita and Mori, 2017).

Since you think that your epistemic group is probably correct about most things within its remit, you will be inclined to trust that group's judgment about those topics. In this respect, trust in the group is analogous to trust in individuals. If you consider your friend Tim to be well informed about literature, and you think his opinion that a novel is worth reading is likely to be true, then you will be inclined to take his recommendation about reading material. The more confident you are in Tim, and the more strongly he recommends a particular title, the more likely it is that you will think it is a good book, even if you have not read it yet.

When people categorize themselves as belonging to a particular group, their confidence in the judgments of other group members solely because of their collective identity becomes more profound. Martin Tanis and Tom Postmes argue that people will trust other identifiable members of their in-group without having to know any other particular characteristics, a courtesy not extended to those outside the group (Tanis and Postmes, 2005). Thus a self-identified Democrat will tend to trust someone simply because they know that person is a fellow Democrat (and will expect reciprocal trust in return), but a Republican will have to personally demonstrate their bona fides before earning trust from a Democrat. One might argue that Tanis and Postme just show that we are inclined to trust fellow epistemic group members, but not the group itself; group membership gives me a reason to trust other members but not the *group*. However, I think the right interpretation is that we view random in-group members less as individuals than as representatives of the group itself, a group we are already prepared to trust. If we didn't antecedently trust the group, then what

possible motivation could there be to trust fellow group members solely on the basis of that membership? Identification with an epistemic group is a powerful force.

There are two related points that deserve acknowledgment, although I lack the space for a full discussion. The first is how people come to learn what the group consensus (or credence level) is about some proposition. Suppose Ivan is a committed member of the Soviet Communist Party. Ivan could learn that the Party believes the best approach to food production is the collectivization of Soviet agriculture in the form of peasant cooperatives by means of an official communiqué from the Party Headquarters, or he could learn of it from his acquaintance Alexsei, another devout and informed Communist. Either way Ivan could update his views about the agriculture issue based on the party line. Obviously adjacent beliefs would come into play as well: Ivan would have to believe that the communiqué was legitimate and not capitalist agitprop, Alexsei was reliably reporting not his own idiosyncratic view but the group judgment, was not poorly informed or dishonest, and so on.

The second point is whether one must self-identify as a member of an epistemic group to be inclined to update one's credences as a result of group collective judgment. The thought is that if you don't think of yourself as a Democrat, then you would not be motivated by Democrats agreeing that P. While that worry has a prima facie appeal, it may be that one could have widespread or pervasive Democratic beliefs, and even be inclined to trust the opinions of the Democratic Party over other political parties, without coming to the conscious realization that one is a Democrat. Finding out that one is part of an epistemic group could be like finding out that one is gay, or good at singing, or is adopted, or likes spicy food, or has large feet—it is a process of self-discovery. In this case one might be an unwitting resident of an echo chamber, like someone who gets all his political views from watching Fox News without the conscious self-reflection of thinking of himself as a conservative.

So far I have argued there is empirical evidence that people tend to trust their own epistemic groups and become more confident as a result of group consensus. None of this should be particularly surprising. More contentious is whether individuals *should* be increasing their confidence in the light of group agreement. There are two reasons to conclude this is a rational thing to do: The Equal Weight View and Aumann's Agreement Theorem.

5. The Equal Weight View

Pettit argues that when an individual disagrees with the majority about propositions that are relatively isolated from other strands in one's web of belief (e.g. whether a car involved in an accident went through a red light or a green one) that deference to the majority opinion is the proper approach. For more complex and psychologically embedded beliefs, he advises that belief revision is warranted only when a supermajority thinks otherwise (Pettit, 2006).

Pettit does not address how an individual should revise her credences when she *agrees* with the group that P is probably true, but simply is less confident that P. The right way to think about this condition is in terms of peer disagreement.

The peer disagreement literature largely assumes three conditions:

Individuals: The disagreeing parties are two individuals; it is not group v. group, or individual v. group.

Peers: The disagreeing individuals are epistemically equivalent; i.e. they are approximately equal with respect to evidential access and reasoning skills.

Competence: The disagreeing peers are both epistemically competent; i.e. not prone to obvious mistakes, not cognitively impaired, not in skeptical scenarios, etc.

David Christensen's well-discussed Restaurant Check Case assumes all three of these elements.

> *The Restaurant Check Case*. Suppose that five of us go out to dinner. It's time to pay the check, so the question we're interested in is how much we each owe. We can all see the bill total clearly, we all agree to give a 20 percent tip, and we further agree to split the whole cost evenly, not worrying over who asked for imported water, or skipped dessert, or drank more of the wine. I do the math in my head and become highly confident that our shares are $43 each. Meanwhile, my friend does the math in her head and becomes highly confident that our shares are $45 each (Christensen, 2007 p. 193).

Christensen argues that since his friend is a competent peer, upon learning of their disagreement over the bill he should be much less confident in his original judgment that each person owes $43. The friend's calculation of $45 provides Christensen a defeater for his own view, and he has no independent reason to discount her opinion apart from their disagreement itself and his assessment that everyone's share is $43. Christensen concludes that he should not give greater credence to his own calculation than to hers. Both he and his interlocutor should revise their self-confidence downward until they reach agreement. If one imagines the other three people at the dinner table witnessing this dispute, they would have no reason to side with Christensen's calculation of $43 or with his friend's calculation of $45 without additional evidence. From the external perspective, both judgments are equally reasonable and they should each be given equal weight. The inside, first-person perspective is not importantly different from the outside, third-person perspective, and so Christensen concludes he should weigh the assessment of his disagreeing interlocutor as equal to his own.

While The Equal Weight View is a popular approach to peer disagreement, it is hardly the only game in town, with others arguing that the first-person perspective does matter due to the ineliminable need for self-trust, or that the second-order evidence of a peer dispute is relevant, but so is the first-order evidence for belief held by each party.[5] I'm not arguing that the Equal Weight View is correct; rather, I am arguing that it is reasonable to hold. There is no received view in the peer disagreement literature about which approach is the obvious front-runner, although Equal Weight is one of the leading contenders. As in other philosophical disputes, one can think that a certain theory is true while also believing that alternative theories have merit and that people of good will and sound faculties might defend them. For example, one could argue that perdurantism is the correct theory of material persistence while also admitting that endurantism has its attractions and it is not unreasonable for someone to defend endurantism. It would be a strident Kantian indeed who thinks that virtue ethicists

[5] (Frances and Matheson, 2019) offer an excellent overview of the salient options.

are flatly irrational in their beliefs.

When two individuals disagree, the Equal Weight View advises substantial revision of their credences; both parties should be far less confident in their own judgments, adjusting their own view to more closely align with the other until they reach convergence. Consider a variant of the restaurant case, one in which four of the diners calculate that the per capita bill is $45 and only one thinks it is $43. In this example, the four in agreement have some reason to reconsider their judgment, but the one outlier has much more reason to update his judgment to conform with theirs. I might think that no one of the other diners is a better mental mathematician than I am, but at the same time I am not better than all of them collectively.

Here's another way to think about the Equal Weight View in the multi-person case. Suppose A disagrees with B, C, D, and E about the restaurant bill, each of whom has come to their own independent judgment about the total. B, C, D, and E agree with each other. Further, A, B, C, D, and E are all equally confident in their judgments, let us say that each person is .8 confident. If we assume that at least one diner correctly calculated the bill, then from the outside view there is a 20% chance that A got it right, a 20% chance B got it right, a 20% chance C is correct, a 20% chance that D is right, and a 20% that E is correct. Since all parties are epistemic peers, Equal Weight recommends taking the outside view. In this case, incorporating the outside view means A should think there's a $(.8 \text{ x } .2) = .16$ probability that he is the one who correctly calculated the bill share. Similar reasoning shows that each of the other four diners should also think there's a .16 chance that they got it right and therefore there is a $(.16 + .16 + .16 + .16) = .64$ probability that at least one of them got the right answer. Since B, C, D, and E all agree, this means they should be .64 confident that the share is $45. Because A now thinks there's only a .16 chance it is $43, A should capitulate and change his mind.

When someone agrees with a group that P is probably true, but is simply less confident that P, this is just another case of one vs. many peer disagreement. They disagree about the likelihood of P. Let's see how this plays out in the restaurant bill case just discussed. Suppose that A, B, C, D, and E all agree that the bill is $45 per person, but A, being a weaker mental mathematician, is only 60% confident in that judgment, whereas each of the others is 80% confident. On the Equal Weight View, A should revise his assessment to more closely conform to the others and increase his credence that $45 is everyone's proper share of the bill.[6]

Even when the disagreement is individual vs. individual, Christensen thinks that revision is the proper approach. His view is increasingly compelling when a lone individual disagrees with a group of epistemic peers who are in agreement with each other.

Echo chambers are reflecting in that group consensus is reflected back to individual members, increasing their confidence. So I far I have argued that there is psychological evidence this actually happens, and argued that the Equal Weight View about peer disagreement helps underwrite its rationality.

---

[6] Compare Begby's similar reflections in defense of his rationalizing account of belief polarization (Begby, 2022 sec. 3).

6. Aumann's Agreement Theorem

Alternative support for the rationality of increasing one's credences to conform with group credence comes from Aumann's Agreement Theorem[7]. Aumann's Theorem states, "If two people have the same priors, and their posteriors for a given event are common knowledge, then these posteriors must be equal. This is so even though they may base their posteriors on quite different information. In brief, people with the same priors cannot agree to disagree" (Aumann, 1976 p. 1236).

Mutual knowledge is when every person in a group of people knows that P. Common knowledge is when every person in a group of people knows that P, everyone knows everyone else knows P, everyone knows everyone knows that everyone knows P, and so on. The distinction between mutual and common knowledge is the crux of Hans Christian Andersen's story of the Emperor's New Clothes. Everyone in town knows that the Emperor is naked, but they are all too fearful to say anything; the result is that no one knows what their neighbor thinks. In other words, there is mutual knowledge that the Emperor is naked, but not common knowledge. When the Emperor struts by at the town parade and the child naïvely blurts out that the Emperor is wearing nothing, then suddenly his nakedness is common knowledge. Everyone in town knows what everyone else knows about the Emperor. Common knowledge is produced when information is spread to individuals via a trust network (Aaronson, 2015).

Aumann assumes that reasoners update their credences by Bayesian reasoning. Suppose that Heloise and Abelard start with the same prior probabilities for all states of the world. Time goes by and they each gain different new information about the world, leading to different posterior probabilities. Let's say that Heloise assigns posterior probability p to a new café opening near campus. Abelard assigns probability q to the same proposition. Heloise and Abelard meet up and during their conversation the possibility of the new café comes up. Heloise tells Abelard that she thinks the chance of it is p, he tells her that he believes the likelihood is q. Notice that neither knows why the other person is giving that estimate, or what the other person's new information might be. That does not matter.

Since it is common knowledge between them that each is a rational Bayesian, and common knowledge what each other's posterior probabilities about the café are, Heloise and Abelard each update their estimates about the café in light of the other's judgment. After the first round, Heloise thinks the chance is p' and Abelard thinks it is q'. Their revised estimates are shared and become new common knowledge between them, leading to further updating to p'' and q'', and so on. So long as the set of possible states is finite, this iterative process will terminate and $p^n = q^n$. Heloise and Abelard will ultimately converge on the same judgment about the probability of the new café, just as they would if they had shared the totality of their evidence with each other (cf. Aaronson, 2005).

Recall that you are in an epistemic group only if you are inclined to trust that group, i.e. only if you think the group is probably right about most of the subjects in its domain. When it becomes common knowledge that the Republicans oppose welfare spending, according to Aumann's agreement theorem, it would be irrational for a Republican to hold out on revising her estimate that welfare spending is bad

[7] Aumann amusingly published this theorem with "some diffidence, since… it is mathematically trivial" (Aumann, 1976). Later he won the Nobel Prize for the proof.

when she sees that everyone else in her trusted group is increasing theirs. Once you notice others in your epistemic group starting to update their credences, you should do the same in the same direction, even if you do not know their evidential basis. When members of your group are increasing their credence in a proposition within the group's domain, the rational thing to do is follow suit. An information cascade—the common knowledge of the posteriors of people with whom you share priors— lifts all epistemic boats.

7. How to get recursion

The final step in building an echo chamber is also the easiest, given the accounts of amplification and reflection above. That is *recursion*.[8] After everyone has updated their confidence via reflection, and are more convinced than ever that P is true, feed that new, higher individual confidence back into the aggregation function. Now the collective group confidence is even higher than it was before, as it again amplifies individual credences. Once again take the new group confidence and reflect it back to individual group members. Again individual confidence gets closer to the group.

For example, suppose your epistemic group is composed of 100 people, each of which is 55.2% convinced that P. Now amplify. By Condorcet's Jury Theorem, collectively your epistemic group is 80% sure that P. After reflection, you revise upwards. Suppose you are now 56.5% sure that P. Everyone else in the group does the same. After another round of amplification the group has 90% confidence that P.[9] Once you see that the group confidence is so high, you use reflection to revise your personal credence upwards again. In the limit, group and individual confidence converge to 100%. This back-and-forth ratcheting upwards is an echo chamber.

8. Objections

One objection is that echo chambers are not truly rational, in that before we can say any account of group attitudes is rational, we need a full Bayesian-style theory of how to aggregate intentional attitudes without being prey to various voter paradoxes and inconsistent beliefs. Furthermore, given List and Pettit's (Arrow-style) Impossibility Theorem that there exists no such aggregation function that satisfies some plausible constraints, there are strong pressures to think that no such theory will be forthcoming (List and Pettit, 2011 p. 51).

However, even if we grant this line of argument, there is a clear sense in which echo chambers are the largely rational and predictable outcome of forces larger than sin-

---

8 (Brugnoli et al., 2019) seems to be alone in recognizing the role of recursion in echo chambers.

9 This calculation uses the formula

$$(P(N,p) = \sum_{k=\left[\frac{N}{2}\right]}^{N} (\,N\,)\,k \times p^k (1-p)^{N-k})$$

where the probability p is initially .5 (a coin toss) + .052 (the probability bias) and N= 100. After the first round of reflection p is stipulated to be .5 + .065. https://demonstrations.wolfram.com/TheoreticalModelCondorcetsJuryTheoremPart1/ provides an interactive calculator for computing this function.

gle agents and their epistemic vices. The mechanisms behind amplification, reflection, and recursion are not due to poor reasoning or cognitive biases. The Condorcet-based logic behind amplification, the increasing of one's credences as a result of group consensus, and the mirroring effect of recursion are all grounded in reasonable epistemic principles.

More strongly, Christensen suggests that fully rational belief cannot require god-like epistemic perfection; if your credences are close enough to the ideal and you lack independent reason for thinking they are too high or too low, then maybe that is good enough to count as totally rational (Christensen, 2009 p. 764). List makes a similar point when he urges that while deference to supermajority testimony may not guarantee complete consistency among one's beliefs, it may nonetheless be a good route to non-blatant, minimal inconsistency, and that is sufficient to count as rational (List, 2014). If human rationality is possible only after we achieve logical or decision-theoretic infallibility, we might as well abandon hope now.

A second objection is that recursion makes believers caught in an echo chamber simply recycle pre-existing information in a way that escalates their confidence, but this is irrational because there is no new evidence being added. They are just working themselves into a tizzy over the same old data when all that data ever supported was (at best) their initial credences. Even if echo chambers really do work through the process of amplification, reflection, and recursion, that just shows that those trapped inside are making an epistemological error.

My response is threefold. First, the preceding objection begs the question in favor of vice theory. Consider again the prisoners' dilemma. One might argue that since both players would be better off if they merely cooperate, mutual defection proves that they are behaving irrationally. While it's true that mutual cooperation would lead to the Pareto-optimal solution, it is rational self-interested behavior that leads to mutual defection, even though that is the worst available solution. The same is true for the examples of Malthusian traps given above. Prisoners' dilemma and similar cases are fascinating and difficult precisely because there is a clear sense in which the participants are driven by reason to make the choices they do. Those choices are what we should logically expect. Yes, believers caught in the closed system of an echo chamber keep increasing their confidence without new ground-level evidence (instead simply consulting and re-consulting the evidence of what their peers think). However, as argued above, they are following quite reasonable principles by doing so and it is wrong to insist they are merely making a mistake. Again, it is what we should logically expect.

Second, Lackey argues that even if many people get their information from the same source and form a belief that P based on that information, their epistemic support for P goes well beyond the source itself. The original information is filtered through their different doxastic frameworks, which brings along potential differences in reliability assessments, belief acceptance, and defeating conditions. Furthermore, each person in the echo chamber is autonomously vouching for the truth of P, which is evidence all by itself. So even if the evidence for ratcheting up credence is entirely dependent on the contents of the echo chamber, "epistemic force is added by each member of the echo chamber through the autonomy of this dependence" (Lackey, 2021a pp. 210-211). If Lackey is right, this is even more reason to view recursion as a legitimate doxastic practice.

Third, recursion is effectively required by Aumann's Agreement Theorem. Suppose Heloise and Abelard start with the same priors and then later (after independent evidence-gathering) arrive at different posterior probabilities. When they re-meet and compare their new posteriors, they cannot rationally let the matter rest there. Abelard must adjust his views in light of the evidence of Heloise's updated assessment, and she must do the same, *mutatis mutandis*. Once they see how much the other has modified their probability assignments, once again they must take each other's new judgments into account as evidence about how much to re-revise their own views. This process is repeated until Abelard and Heloise converge on the same posteriors. It is when we add amplification and reflection to this process that we get an upward-spiraling drive to certainty in an echo chamber.

9. Benefits of the present analysis

There are numerous benefits from the present analysis. It is prima facie puzzling how people can be completely convinced in the truth of total nonsense (e.g. Qanon), but the present account of echo chambers serves to explain it. Take Carnap's example of metaphysical meaninglessness: "The Nothing is prior to the Not and the Negation… Anxiety reveals the Nothing" (Carnap, 1959). If Carnap is right and that is a pseudo-statement, then it is not even false. Nonetheless, Condorcet can build agreement, which is why the Jury Theorem is not fundamentally about truth values, or about convergence to the truth, at all. Rather, it is convergence to certainty. Imagine a Heideggerian echo chamber. Suppose that each member of the group thinks that it is 60% likely that "The Nothing is prior to the Not and the Negation… Anxiety reveals the Nothing." Per Condorcet, the group assessment will be higher than 60%. The higher group credence is reflected back to the individual group members who then update their judgments to more closely reflect their trust in the group's collective wisdom. Those higher credences are again aggregated and the group's credence increases; recursion then escalates it to certainty.

One might object that few Heideggerians or anyone else in an echo chamber really assign probability one to their beliefs. The extent to which that is true is an empirical matter, but surely there will be dropouts and defectors who, for whatever reason, leave an epistemic group. In addition we should not expect the real world to have perfect information transfer between individuals and groups. So not everyone is going to wind up with absolute confidence in their beliefs as the result of an echo chamber. The point is that as long as one belongs to an epistemic group and there is the recursive amplification and reflection of beliefs between the individual and the group, there will be rising pressure to increase one's credence.

A second benefit is that the present account explains why echo chambers produce robust epistemic insularity. Many writers on echo chambers either equate them with homogenous and polarized communities (Brugnoli et al., 2019), highlight that they insulate those within from rebuttal (Jamieson and Cappella, 2010), or point out that they reinforce a set of views and discourage dissent (Boyd, 2019). Nguyen distinguishes echo chambers from epistemic bubbles. He argues that bubbles are frail and easy to pop—an epistemic bubble is merely due to a limited diet of information, and excludes other resources through omission. Such an oversight could happen accidentally, and "simple exposure to excluded voices can relieve the problem" (Nguyen,

2020 p. 153). Echo chambers he characterizes as actively discrediting outsiders, so that mere exposure to outside or contravening arguments will have no effect on those inside the chamber.

Nguyen is right to distinguish bubbles and echo chambers. He is also right that *some* epistemic bubbles are fragile and may be the accidental result of an information desert. However, not all bubbles are like that. In particular, on the account I have given, we should expect echo chambers to be self-sealing, and to produce robust epistemic bubbles that are quite difficult to burst. Here is why. Suppose that you are mildly convinced that politician P has your best interests at heart and deserves your vote. You think it is somewhat better than a coin toss, anyway. If you are just 51% confident in that politician, you will be open to hearing counterarguments, considering other candidates, willing to entertain evidence that the politician is corrupt or dishonest, and so forth. After all, you are only slightly in favor of that politician, and acknowledge that you could easily be mistaken. Your vote is not written in stone. However, if you are 99.99% sure that the politician is the best candidate you will do none of those things. Why bother to look at other evidence if you are 99.99% confident?

Lackey writes, "if I learn about climate change from a reputable environmental scientist, there is only the danger of acquiring false beliefs in also consulting a climate change denier" (Lackey, 2021a p. 216). Lackey is completely convinced in the truth of anthropocentric climate change, and so looking at contravening information is either a pointless waste of time or would harmfully push her away from the truth. Hers is a completely reasonable and common attitude. However, being in an echo chamber will take your 51% confidence in politician P and turn it all the way up to certainty. When you are certain, or nearly so, there is no reason to look at other perspectives; competing sources of information/ data/ arguments are irrelevant at best and dangerous at worst. Lackey believes she is in an echo chamber, just a good one. Of course, everyone thinks their echo chamber is virtuous—that is just how they work. Robust epistemic bubbles are not a component of echo chambers, as often thought, but a predictable effect of them.

J. Adam Carter and Fernando Broncano-Berrocal warn that "Echo-chambers [are] where the available arguments and information lean in the same direction, potentially leading to group polarization."[10] A third benefit of the present account is that it easily explains how that happens and correctly predicts that polarization will get worse as intra-group communication improves. When any like-minded collection of people forms an epistemic community and share their beliefs with each other, amplification will ensure that the group collectively holds such beliefs more strongly than the average member. Those beliefs could be true, false, or even nonsense, but once the echoing mechanism gets started, everyone in it will become increasingly convinced. When it comes to topics that are political, religious, or otherwise unsettled and controversial, we should expect a diversity of epistemic groups: those who believe that transsexual women are really women and those who believe they are deluded men, those who maintain racism is a serious problem and those who aver that it is overblown, those who advocate a massive redistribution of wealth and those who think taxation is theft. All will eventually be completely certain of their views within

[10] (Carter and Broncano-Berrocal, 2021 p. 14) and (Boyd, 2019 p. 63) concurs: "Increased confidence within epistemic bubbles and echo chambers creates another oft-discussed epistemically pernicious effect, namely *group polarization*." [italics in original]. Also see (Begby, 2022).

an echo chamber.

In the past, it was more difficult for those with fringe beliefs—the earth is flat, the Illuminati secretly run the government, the Sandy Hook mass shooting was a hoax—to effectively form an epistemic group because they were relatively few and far between. With the internet and social media, it is now easy to find your fellow conspiracy theorists and exchange ideas. In addition, it is far easier to silo your information feed when so much of it is conducted in virtual interactions. In their empirical study of a large Facebook dataset, Emanuele Brugnoli and colleagues find "that polarized users reinforce their preexisting beliefs by leveraging the activity of their like-minded neighbors. Such a trend grows with the user engagement" (Brugnoli et al., 2019). On my view of echo chambers, they are an automatic outcome of amplification and reflection, and the acceleration to maximum credence gets faster when recursion speeds up. This is just what Brugnoli et al. find: social media allows for quicker feedback from one's epistemic group as well as allowing faster group aggregations of beliefs. Because social media accelerates recursion in an echo chamber, and there is a diversity of echo chambers across a range of opinions, it drives swiftly towards a more pronounced polarization of beliefs.

The main alternative account of echo chambers in the literature is Vice Theory: echo chambers exist because we screwed up. Brugnoli et al. claim they arise due to confirmation bias and a refusal to admit error, and Boyd writes that echo chambers are a kind of "epistemically pernicious group" (Boyd, 2019 p. 61). Nguyen maintains that echo chambers turn individually virtuous epistemic practices into collective epistemic vice, calling them perversions of trust. He likens them to cults and suggests they are intentionally set up by malicious actors (Nguyen, 2020 pp. 147, 149, 155). Avnur writes that "echo chambering is bad" (Avnur, 2020 p. 579) and it is bad because it is the result of cognitively biased inquiry (p. 580) that "make[s] a familiar problem, motivated reasoning, much worse" (p. 590).

Reaching for virtue epistemic tools is an expected and familiar thing to do. The idea that we need to cultivate open-minded communities of epistemic trust, and our failure to do so is the vice that results in echo chambers is attractively self-critical. However, there are two problems with Vice Theory. First, not all echo chambers are vicious, nor is everyone within an echo chamber unfairly prejudiced against external views.[11] Vice Theory struggles to accommodate this fact. Second, Vice Theory lacks the explanatory power of the model defended here.

I have already argued that echo chambers arise from fundamentally rational forces in group reasoning. Vice Theory fails to acknowledge that not all echo chambers amplify falsehoods among their members. Nguyen claims that echo chambers are a reverse-Mandevillian intelligence that turns generally good epistemic practices into locally unreliable ones (Nguyen, 2020 p. 156). Those trapped in such chambers wrongly discredit and distrust the views of outsiders while embracing the epistemic credentials of those inside. While many of the familiar echo chambers that come naturally to mind (socially isolated religious sects, 9/11 truthers, Fox News, etc. (Licari, 2020) may fit Nguyen's description, there are also many—like Lackey's community of climate scientists—that do not. If an echo chamber starts with group members who have high individual credences in the truth, it will virtuously increase those credences. Lackey has a stronger grip on the truth of climate change by systematically discrediting

[11] (Elzinga, 2022) also recognizes this.

and ignoring the deniers.

If you are nearly certain that a proposition is true, and it is certified by your trusted epistemic peer group, it is proper to disregard evidence to the contrary. There is no value in taking seriously sophistical arguments or spurious data. As (Fantl, 2021 p. 646, cf. p. 654) observes, "we should want to inhabit truth-conducive echo chambers." Vice Theory gets this wrong: it holds that all echo chambers are due to cognitive bias or poor epistemic hygiene and none can be virtuous. Even if all Lackey's sources on climate change are either the scientific literature or professional climate researchers, that only shows she is caught in a scientific echo chamber. Since none are virtuous, she will still need to respectfully assess the arguments of deniers to escape. As she is quick to note, that is surely the wrong conclusion.[12]

Even a more modest vice theorist who thinks echo chambers are generally epistemically pernicious would then have to find a way to distinguish between reasonably dismissing absurdly implausible alternative views (when you are in a virtuous chamber) and unreasonably dismissing them (when you are in a vicious one). The problem for the vice theorist now is the structure of the two echo chambers is the same, and given their epistemic resources and conviction in their beliefs, both groups use the same reasoning to spurn external evidence. Short of assuming a perspective *sub specie aeternitatis* to validate one echo chamber over the other, the vice theorist is at a loss to explain why one group is acting badly by disregarding outsiders and the other is not.

Vice Theory does not predict echo chambers becoming more prominent or polarization becoming more extreme. Presumably a vice theorist would have to explain post hoc why people are increasingly failing to cultivate epistemic virtues, and that more distrust, arrogance, insularity, and poor data assessment leading to polarization are the results. However a vice theorist might develop that argument, it would be a response to the observation of increased polarization. There is no particular a priori reason to expect increased polarization if echo chambers are due to vice. My proposal, on the other hand, does predict it. As argued earlier, the faster information about group credence is reflected back to individuals, the quicker they can increase their own confidence. The internet and social media thus speed up recursion and any tendency to polarize is accentuated. Vice Theory is right that there is a relationship between echo chambers and epistemic vice, but mischaracterizes the connection. It is the feedback loop of echo chambers that start with falsehoods that gives rise to vice.

Echo chambers are not a bug in group reasoning, but a feature. Viewing them as the result of vice misses much that makes them interesting. Echo chambers are practically automatic once certain conditions are satisfied; they are not a failure of any kind, but rather a completely predictable, rational, logically expected outcome. That doesn't mean they lead to good results; phantom traffic jams arise even when all drivers are doing the best they can and following the rules of the road. Echo chambers gone bad are a kind of logical traffic jam, where you can start mildly persuaded in dubious claims, and proceed step by reasonable step to total conviction in falsehoods.

In sum, echo chambers are the outcome of fundamentally rational forces in group reasoning: amplification, reflection, and recursion. The computer programming cliché of "garbage in, garbage out" applies here as well—when an echo chamber gets

[12] Which is partly why Lackey rejects vice theory. She argues the problem with echo chambers is not their structure, but (when containing falsehoods) their content (Cf. Lackey, 2021a).

started with bad information or even abject nonsense, it will build in intensity until its members are walled off from the truth. It is tempting to see those trapped in such circumstances as stupid or wicked, but that is a mistake. Thus the account defended here respects the principle of charity. Furthermore, groupthink, robust epistemic bubbles, and increasing polarization are all predicted and explained by the Reasonable Expectations model of echo chambers in a way that Vice Theory struggles to achieve.

Echo chambers are not inherently pernicious, and we would all do well to be ensconced in ones that promote the truth. Certainly many echo chambers are bad ones, but we should not expect a general solution to the bad ones any more than we should expect a general solution to prisoner's dilemmas or other problems of suboptimal equilibria. All we are going to get are strategies and piecemeal approaches. One take-away is that we ought to be cautious about which epistemic groups we join. Epistemic groups are not the same as echo chambers, since we might pay no attention to the collective judgment of groups to which we belong and thereby not get caught up in the feedback loop described here. Nonetheless, belonging to the wrong epistemic groups runs the risk of downward spiraling echo chambers. While the right ones may still generate echo chambers, at least they will be ones that reinforce the truth. The problem is how to make sure we are in the right echo chambers, when it is an analytic truth that everyone already thinks their own epistemic group is likely to be correct.[13]

Acknowledgements

Thanks to audiences at Clemson University, Universidad de la Laguna, New Mexico State University (especially Mark Walker), the XXV World Congress of Philosophy, and Luke Semrau for criticisms of previous versions of this paper.

References

Aaronson, Scott (2005), 'The Complexity of Agreement', Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing 634-43.

——— (2015), 'Common Knowledge and Aumann's Agreement Theorem', *Shtetl-Optimized*, https://www.scottaaronson.com/blog/?p=2410.

Asch, Solomon E. (1955), 'Opinions and Social Pressure', *Scientific American*, 193 (5), 31-35.

Aumann, Robert J. (1976), 'Agreeing to Disagree', *The Annals of Statistics*, 4 (6), 1236-39.

Avnur, Yuval (2020), 'What's Wrong With the Online Echo Chamber: A Motivated Reasoning Account', *Journal of Applied Philosophy*, 37 (4), 578-93.

Baumgaertner, Bert (2014), 'Yes, no, maybe so: a veritistic approach to echo chambers using a trichotomous belief model', *Synthese*, 191, 2549-69.

Begby, Endre (2024), 'From Belief Polarization to Echo Chambers: A Rationalizing Account', *Episteme*, 21 (2), 519-539.

[13] Ethical considerations: All research and writing is my own, and there are no other ethical principles relevant to the paper.

Bond, Rod and Peter B. Smith (1996), 'Culture and Conformity: A Meta-Analysis of Studies Using Asch's Line Judgment Task', *Psychological Bulletin*, 119 (1), 111-37.

Boyd, Kenneth (2019), 'Epistemically Pernicious Groups and the Groupstrapping Problem', *Social Epistemology*, 33 (1), 61-73.

Brugnoli, Emanuele, et al. (2019), 'Recursive Patterns in Online Echo Chambers', *Scientific Reports*, 9

Carnap, Rudolf (1959), 'The Elimination of Metaphysics Through the Logical Analysis of Language', in Ayer, A. J. (ed.), *Logical Positivism* (New York: The Free Press),

Carter, J. Adam and Fernando Broncano-Berrocal (2021), *The Philosophy of Group Polarization*, (London: Routledge).

Christensen, David (2007), 'Epistemology of Disagreement: The Good News', *Philosophical Review*, 116 187-218.

——— (2009), 'Disagreement as Evidence: The Epistemology of Controversy', *Philosophy Compass*, 4/5 756-67.

Dietrich, Franz and Kai Spiekermann (2020), 'Jury Theorems', in Fricker, Miranda, et al. (eds.), *The Routledge Handbook of Social Epistemology* (New York: Routledge), 386-96.

Elzinga, Benjamin (2022), 'Echo Chambers and Audio Signal Processing', *Episteme*, 19 (3), 373-93.

Estlund, David (2008), *Democratic Authority: A Philosophical Framework*, (Princeton: Princeton University Press).

Fantl, Jeremy (2021), 'Fake News vs. Echo Chambers', *Social Epistemology*, 35 (6), 645-59.

Flynn, Morris R., et al. (2023), 'Traffic Modeling— Phantom Traffic Jams and Traveling Jamitons', *https://math.mit.edu/traffic/*,

Frances, Bryan and Jonathan Matheson (2019), 'Disagreement', in Zalta, Edward N. (ed.), *The Stanford Encyclopedia of Philosophy* (https://plato.stanford.edu/archives/win2019/entries/disagreement/.

Fujita, Yuki and Kazuo Mori (2017), 'Group versus Individual Reward in the Asch Experiment without Confederates', *Open Journal of Social Sciences*, 5 396-402.

Jackson, Elizabeth G. (2020), 'The Relationship Between Belief and Credence', *Philosophy Compass*, 15 (6), 1-13.

Jamieson, Kathleen Hall and Joseph N. Cappella (2010), *Echo Chamber: Rush Limbaugh and the Conservative Media*, (Oxford: Oxford University Press).

Lackey, Jennifer (2021a), 'Echo Chambers, Fake News, and Social Epistemology', in Bernecker, Sven, Amy K. Floweree, and Thomas Grundmann (eds.), *The Epistemology of Fake News* (Oxford: Oxford University Press), 206-27.

——— (2021b), *The Epistemology of Groups*, (Oxford: Oxford University Press).

Licari, Peter R. (2020), 'Sharp As a Fox: Are foxnews.com Visitors Less Politically Knowledgeable?', *American Politics Research*, 48 (6), 792-806.

List, Christian and Philip Pettit (2011), *Group Agency: The Possibility, Design, and Status of Corporate Agents*, (Oxford: Oxford University Press).

List, Christian (2014), 'When to Defer to Supermajority Testimony—and When Not', in Lackey, Jennifer (ed.), *Essays in Collective Epistemology* (Oxford: Oxford University Press), 240-49.

Nguyen, C. Thi (2020), 'Echo Chambers and Epistemic Bubbles', *Episteme*, 17 (2), 141-61.

Pettit, Philip (2006), 'No Testimonial Route to Consensus', *Episteme*, 3 (3), 156-65.

Ranalli, Christopher and Finlay Malcolm (2023), 'What's So Bad About Echo Chambers?', *Inquiry*, forthcoming

Sunstein, Cass R. (2006), *Infotopia: How Many Minds Produce Knowledge*, (Oxford: Oxford University Press).

Tanis, Martin and Tom Postmes (2005), 'A Social Identity Approach to Trust: Interpersonal Perception, Group Membership, and Trusting Behavior', *European Journal of Social Psychology*, 35 413-24.

Yudkowsky, Eliezer (2017), *Inadequate Equilibria: Where and How Civilizations Get Stuck*, (Berkeley, CA: Machine Intelligence Research Institute).