

Epistemically Pernicious Groups and the Groupstrapping Problem

Forthcoming in *Social Epistemology*

Recently, there has been growing concern that increased partisanship in news sources, as well as new ways in which people acquire information, has led to a proliferation of *epistemic bubbles* and *echo chambers*. In an *epistemic bubble* one acquires information from a limited range of sources, ones that generally support the kinds of beliefs that one already has, perhaps as a result of behind-the-scenes digital curation. *Echo chambers* function in the same way, but possess the additional characteristic that certain beliefs are actively reinforced, sometimes in the form of portraying outside sources as being misleading or untrustworthy. The existence of epistemic bubbles and echo chambers are often taken to explain the apparent explosion of “fake news”, as well as other contemporary epistemic maladies, such as group polarization. Despite representing some of the most crucial and practical epistemic problems we face today, serious philosophical analyses of epistemic bubbles and echo chambers are still nascent: there are, for example, open questions about how epistemic bubbles and echo chambers form, how they persist, and why it seems that they are so difficult to break out of.

Here I will first argue that we should conceive of epistemic bubbles and echo chambers as types of *epistemically pernicious groups*. As I outline below, there are two primary problems created by such groups: first, members of said groups will tend to have an unwarranted degree of confidence in their beliefs, and second, these groups encourage the stagnation of processes of inquiry within its members. I argue next that while analyses of epistemically pernicious groups have typically focussed on relationships between individuals, one of the most significant problems with such groups is a problem of *group testimony*: in an epistemically pernicious group the group itself is treated as a source of information by a member, resulting in increased

confidence in member belief, which in turn results in attributing an unwarranted degree of credibility to the group itself on the basis of increased member confidence. I call this process *groupstrapping*. I argue that by recognizing groupstrapping as an illicit method of forming and updating beliefs we can make progress on some of the open questions concerning epistemically pernicious groups.

The paper will proceed as follows: in section 1, I outline some of the recent discussions of epistemic bubbles and echo chambers, with a focus on their negative epistemic consequences as a product of relationships between individuals. In section 2, I argue that we should conceive of epistemic bubbles and echo chambers as epistemically pernicious groups, and show how such groups themselves are treated as a source of testimony by members within said groups. In section 3, I describe groupstrapping as an illicit method of increasing confidence in beliefs of individuals and groups themselves. In section 4, I argue that recognizing the groupstrapping problem can help us make better sense of the problems that epistemically pernicious groups create, and distinguish the groupstrapping problem from related problems pertaining to rational belief formation and updating. In section 5, I conclude with a brief suggestion about how we might think about strategies for breaking out of epistemically pernicious groups.

1. Epistemic Bubbles and Echo Chambers

C. Thi Nguyen (2018) defines an epistemic bubble as “an informational network from which relevant voices have been excluded by omission”, where such exclusion can be intentional – for example, in cases in which others are excluded out of bias – or unintentional – for example, in social media groups that tend to be populated by those who agree with us, not because we

actively seek them out, but simply because people with similar views tend to flock together. Eli Pariser (2011) describes the same phenomenon using the term “filter bubble” to refer specifically to online environments in which one tends to be exposed to information that conforms to one’s pre-existing views, environments that are created by algorithms that monitor browsing behaviors¹. As Pariser argues, epistemic bubbles are epistemically pernicious in two primary ways:

First, the filter bubble surrounds us with ideas with which we’re already familiar (and already agree), making us overconfident in our mental frameworks. Second, it removes from our environment some of the key prompts that make us want to learn. (84)

The first problem, that of overconfidence, stems from a failure to properly update one’s confidence in a belief in response to the available evidence: the mere fact that many others agree with you does not warrant an increase in confidence in your belief if everyone is getting their evidence from the same sources and reaching the same conclusions in the same way. The second problem pertains to an epistemic bubble’s tendency to make us complacent: the “key prompts” that Pariser refers to are instances in which one comes across new ideas from diverse sources, ones that we might not necessarily end up agreeing with, but which at least get us to reflect upon and scrutinize the beliefs that we already hold. Without these prompts, one’s processes of inquiry stagnate, as one will tend to be content with holding onto beliefs that have gone unchallenged, and because one cuts oneself off from new and potentially better sources of information.

Echo chambers are similar to epistemic bubbles insofar as they, too, exclude sources of information. Unlike epistemic bubbles, however, echo chambers also actively reinforce a

¹ For my purposes here I take “epistemic bubble” and “filter bubble” to be synonymous.

particular belief or set of beliefs (hence one's views are "echoed back"). Recent discussions of echo chambers outline different ways in which this reinforcement can take place:

[By "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...[it] creates a common frame of reference and positive feedback loops for those who listen to, read, and watch those media outlets...We mean 'echo' is a second sense as well: each [media] outlet legitimizes the other. (Jamieson and Cappella, 2009: 76)

An echo chamber is a sociological setting where peoples' prior beliefs are 'echoed back' giving the impression that their beliefs are correct. Examples of echo chambers include media networks...blogs and online niche forums...and even the right setting in a classroom of peers...Such sociological mechanisms systematically foster confirmation bias and belief amplification, resulting in significantly high levels of agreement between individuals. (Baumgaertner, 2014: 2549-50)

[Echo chambers] are ideologically congruent and [homogeneous] environments in which political views are not debated but instead reinforced and amplified thus paving the way to increased polarization. (Justwan et al., 2018: 2)

[T]o the extent that individuals expose themselves to information that simply reinforces their existing views...greater access to information may foster selective exposure to ideologically congenial content, resulting in an "echo chamber" environment that could facilitate social extremism and political polarization. (Barbera et al., 2015: 1531)

While epistemic bubbles merely *exclude* certain views, echo chambers *reinforce* a set of views.

From the discussions above we can see a number of ways in which this can take place: by simply

having a set of views repeated, by discouraging dissent with the group, and/or by discounting opposing views or sources. We can thus identify two of the major epistemically pernicious characteristics of echo chambers. The first is the same as that present in epistemic bubbles, namely that in an echo chamber individuals tend to increase their confidence in their beliefs to an unwarranted degree. Unlike epistemic bubbles, however, this effect is produced in two ways: one from agreement with others within the echo chamber, and another from active reinforcement. As Nguyen (2018) argues, it thus becomes very difficult to change the mind of someone within an echo chamber: it is often the case that outside information, rather than prompting reflection, is treated as false or misleading, and thus the presentation of information that suggests that my belief is mistaken results in an increase in confidence in said belief, rather than a decrease. Indeed, one may be warranted in increasing one's confidence in one's belief if a source known to be untrustworthy conveys information that indicates that one's belief is false. However, this method of updating one's belief becomes pernicious when dissenters are considered untrustworthy *simply because* they are dissenting. This is one of Nguyen's primary concerns with echo chambers: echo chambers will often present the mere fact of disagreement to be an indication of untrustworthiness.

Increased confidence within epistemic bubbles and echo chambers creates another oft-discussed epistemically pernicious effect, namely *group polarization*, the phenomenon in which individuals within a group tend to believe or disbelieve the same things with very high degrees of confidence. Cass Sunstein (2017) argues that there are three main factors that contribute to polarization. First, there are limited argument pools: when everyone is on roughly the same side, everyone appeals to the same arguments to support their beliefs, making it seem as though said beliefs are better supported than they actually are. Second, individuals are driven by reputational

considerations: one wants to be viewed favorably by others, so one is wont to shift one's views towards the more extreme in order to be accepted. Finally, there are the aforementioned effects of corroboration, namely that one will be unduly confident in one's views because of widespread agreement, agreement that stems from shared evidence and arguments.

There are, then, two primary types of epistemic problems created by epistemic bubbles and echo chambers: unwarranted levels of confidence in members' beliefs, often leading to group polarization, and a stagnation of processes of inquiry. In the above discussions, these two problems have been framed in terms of the effects of individuals upon other individuals within epistemic bubbles and echo chambers: it is the fact that others individuals within my group agree with me that result in my increased confidence in my beliefs, and it is my reliance on these individuals that causes my belief-forming processes to stagnate. There is, however, another way that these problems arise, namely as the result of the interaction of members of a group and the group itself.

2. Epistemically Pernicious Groups and Group Testimony

Here I want to establish three claims: first, that epistemic bubbles and echo chambers should be conceived of as groups; second, that groups are often legitimately treated as sources of testimony; and third, that what a group is warranted in testifying is a function of the warrant of its members.

Why should we conceive of epistemic bubbles and echo chambers as groups? The above discussions tend to talk in terms of *environments* (or sometimes "media spaces" or "sociological settings"), specifically environments in which one acquires information. It is not clear, however,

what constitutes an environment. Social media are often taken to be examples: one's Facebook page or Twitter feed might constitute environments, but so might the curated internet experience that has been generated by one's browsing behavior, or perhaps the news channels that one watches on television, or the newspapers that one chooses to read, etc. These environments, however, are shared with others, namely those with similar browsing habits, and who watch and read the same things. As a result, there is reason to think that by sharing the same sources of information, and participating in the same communities online or offline, one becomes a member of a *group*, namely a group that shares information that promotes certain types of views².

To get a better sense of how epistemic bubbles and echo chambers function as groups we can briefly consider some characteristics of groups in general. Kathleen Ritchie (2013), for example, argues that groups should be conceived of as “realizations of structures”, where a structure is constituted by nodes that are functionally related to one another, and in which membership in a group consists in occupying one such node. A baseball team, for example, constitutes a group on this model, since members of the team are related to one another according to the functions that they serve, none occupying any more important of a position than any other; a corporation, on the other hand, can similarly constitute a group, although its structure will be such that different members occupy more important positions according to the way the corporation is organized. Given these criteria, we can conceive of epistemic bubbles and echo chambers generally as types of groups in which information conforming to a certain

² One of my primary aims in changing the discussion from environments to groups is for the sake of precision: in none of the above discussions is the concept of environment defined in detail. Furthermore, while we often treat groups as sources of information (as I discuss below), it is not clear in what sense an *environment* can serve as a source of information. If one takes there to be no difference between an environment and a group then this change may be interpreted as merely terminological.

mandate or viewpoint is acquired and shared, and in which trusted sources or perceived experts may occupy positions of higher importance.

These groups will be epistemically pernicious insofar as they are ones in which the belief-forming and updating methods employed by members of such groups will tend to lead to said members having false and/or unwarranted beliefs. Not all groups that have the mandate of disseminating and sharing certain types of information, then, will necessarily be epistemically pernicious. For example, one can be a member of the Facebook high school reunion group “Class of ‘99” in which a particular type of information is being shared (namely, information pertaining to the goings on of the members of the class of ’99) without exemplifying any of the aforementioned epistemically pernicious characteristics: one may or may not think that the class of ’99 was a particularly strong one, for example, or that the high school in question was a particularly good one, and still be a member of the group, and there is no requirement that one distrust members of any other high school reunion groups. Furthermore, promoting a particular mandate does not necessarily make a group epistemically pernicious. For example, to be a member of the “Persian Cat Aficionado” group will require that one share the mandate of the group that Persian cats are an excellent choice of pets. One is not, however, in virtue of being a member of that group, necessarily cut off from information from any other sources, nor does one need to be distrustful in any way of competing views (one could very well also be a member of the “Siamese Cat Aficionado” group, for example, without risk of hypocrisy). Finally, a group that is dismissive of the views of outside individuals and groups, but does not endorse any particular positive mandate, will not necessarily be epistemically pernicious, either. Members of the society of skeptics, for example, will generally be distrustful of any views expressed by

anyone outside of the group, but this itself will not necessarily lead to unwarranted beliefs or a stagnation of inquiry.

By conceiving of epistemic bubbles and echo chambers as types of groups, we can draw attention to an additional way in which the epistemically pernicious aspects of said groups come about, namely as a result of *group testimony*. While historically only individuals have been considered to be sources of testimony, social epistemologists have recently made the case that we should include groups as potential testifiers, as well. There is good reason to include groups in the category of sources of testimony, as groups like government organizations, laboratories, universities, clubs, online communities, etc., are commonly appealed to as sources of information. As is the case for individuals, some groups are better sources of testimony than others; similarly, questions pertaining to the conditions necessary for one to acquire knowledge on the basis of testimony from individuals also apply equally well to groups (for example, we can ask whether a recipient of testimony needs to trust the testifier in an explicit or implicit way, whether one needs to possess positive reasons or lack negative reasons to think that the testifier knows what they are talking about, etc.).

I will not defend any specific theory of testimony here; I only rely on the claims that groups can be sources of testimony, and that there is nothing in principle illicit about forming and updating one's beliefs on the basis of group testimony. The question that is of greater concern here is how to determine what a group is warranted in testifying. It is a matter of debate as to how we should conceive of the epistemic status of a group. For instance, although we certainly talk as if groups believe things – for example, sentences like “it is NASA's belief that a hundred more exoplanets will be discovered this year”, and “it is the laboratory's belief that an effective treatment will be devised soon” seem perfectly felicitous – it is not clear how we should

make sense of these claims, as it is not clear whether groups are the kinds of things that can have beliefs³. Without taking a stance on this question, I think it is more natural to say that a group can be *warranted in believing* something. To say that a group has an epistemic position with respect to some proposition as I will consider it here, then, is to make a claim about the nature of a group's propositional, as opposed to doxastic, warrant.

What does seem clear is that what a group is warranted in believing will be a function, in some way, of the warrant of the members within that group. There are in general two different ways we can think about the epistemic relationship between a group and its members: a *deflationary* sense and an *inflationary* sense. As Debra Tollefsen (2007) argues, on the deflationary account a group is warranted in testifying a proposition just in case the majority of its members are; whereas on an inflationary account a group may be warranted in testifying a proposition that few or none of its members are. Tollefsen prefers an inflationary account, for the reason that there are many groups that are comprised of members who contribute towards a common epistemic goal. As a result, there may be only a few members at the hierarchical top, or else a designated spokesperson of the group who have beliefs that express the collective work of its members, beliefs that the majority of the members do not actually themselves hold (we can think, for example, of a laboratory comprised of members who all work on separate parts of a larger project).

Debates about whether we should accept a deflationary or inflationary account of group warrant parallel discussions about how we should conceive of the metaphysical status of groups in general. For instance, one might hold a view in which groups are nothing over and above a

³ For a review of recent discussions of group belief, see Lackey (2016).

collection of individuals, and thus believe reference to “group beliefs” and “group testimony” should be considered metaphorical. On the other hand, one might hold a view in which groups are entities in their own right, ones that are constituted by individuals but have a distinct ontological status. Consider, for example, how Philip Pettit (2010) distinguishes *aggregates* from *integrations*. According to Pettit, aggregates are “groups that do not reason at all or that do not impose the discipline of reason at the collective level”, while integrations are purposeful, and reason at the level of the collection as opposed to the level of the individual (178). One way that a group qua integration can reason is through a decision-making procedure, in which the beliefs of its members are weighed according to some relevant set of factors (one of Pettit’s central examples involves unions making decisions on the basis of votes of their members). When a group constitutes an integration, then, beliefs attributed to the group can differ from those held by its members, in some cases to the extent that the groups’ belief may not be shared by the majority of its members.

How should we conceive of epistemically pernicious groups? As discussed above, some epistemically pernicious groups have a mandate in terms of the dissemination of certain types of information, and, perhaps especially in echo chambers, such groups express views that appear to not be mere consolidations of the views of its members, especially when group views are expressed in such a way as to minimize dissention. There also do appear to be epistemically pernicious groups that are not organized in such a way that would qualify them as an integration in something like Pettit’s sense: for example, Nguyen describes cases of social media groups that consist of like-minded individuals, but ones that do not communicate with each other in such a way that involves any specific organization structure or decision-making procedure. Such groups would perhaps be best conceived of as mere aggregates.

I will not here make a claim about whether all epistemically pernicious groups should be conceived of as aggregates or integrations. For my purposes, I take there to be good reason to think that such groups come in both forms. Regardless, then, of whether we prefer a deflationary or inflationary account of group testimony, and regardless of whether we conceive of groups as mere collections or as integrations, we can nevertheless establish a general relationship between the warrant of a group and the warrant of its members. For instance, the more epistemically well-positioned the members of a group are towards a proposition and/or reasons that support a belief in that proposition, the better epistemically well-positioned the group will be. As such, if every member of a group has a weak epistemic relationship with p , then it will be likely that the group itself also has a weak epistemic relationship with p . For example, it is the position of the flat-earth society that the earth is flat, because each of its members believe that the earth is flat. Furthermore, the flat-earth society is not epistemically well-positioned towards this proposition because none of its members are warranted in believing that the earth is, in fact, flat. On the other hand, a laboratory team that is comprised of scientists working on their own individual parts of a larger problem – whether it is the case that p – will be epistemically well-positioned towards p so long as its members are epistemically well-positioned towards reasons that support a belief that p . While the majority of the lab members will not have any epistemic relationship with p itself, then, each scientist's contribution involves more or less warranted beliefs, which will make the group itself more or less warranted in its belief that p . Again, depending on how group warrant is determined, and whether a group is an aggregate or an integration, this relationship will be more or less proportional. Nevertheless, it remains that increased member warrant will end to increase group warrant.

Just as individuals can be confident in their beliefs to a degree that is not properly warranted by their evidence, so too can groups be ascribed a degree of confidence that outruns their properly ascribed warrant. Again, there are metaphysical questions pertaining to whether groups are the kinds of things that can be confident. Nevertheless, it seems that just as we can ascribe warrant to a group, so too can we ascribe confidence: the flat-earth society, for example, is very confident that the earth is flat, whereas the skeptic's society is not very confident in anything. As was the case with group warrant, group confidence will be a function of member confidence, however one wants to construe it. In general, one way that group confidence will outrun group warrant, then, is if member confidence generally outruns member warrant.

The problem that I want to draw attention to as it pertains to epistemically pernicious groups occurs when groups are ascribed a degree of confidence that outruns their warrant, and are thus perceived to be more credible sources of information than they actually are. To generate this problem we need one more bit of terminology. There are two types of groups that we treat as testimonial sources: groups that we are not members of, and groups that we are members of. Call this latter type of group a *self-membership group*. Relying on the testimony of self-membership groups is again a common occurrence: for example, I am a member of and receive testimony from my university, from my department, from my social media groups, etc. While there is not, in principle, anything illicit about the practice of relying on the testimony of self-membership groups, such reliance can lead to groups developing the epistemically pernicious characters. I describe this process next.

3. Groupstrapping

Here are two key claims that I want to emphasize from the discussion so far:

- 1) The extent to which a group is warranted in believing a proposition is (at least in part) a function of the relevant warrants of its members⁴.
- 2) Groups (specifically, self-membership groups) can be sources of testimony.

Given these claims we can generate the problem of how a relationship between a group and its members can lead to that group becoming epistemically pernicious, what I am calling the groupstrapping problem. Consider a group G , and a member of that group, m . According to (2), groups can be sources of testimony (and, as we have seen, seeking out information from self-membership groups is not in itself illicit), so we can consider a case in which m seeks out information from G , and G testifies that p to m . On the basis of G 's testimony, m then either comes to form a new belief that p , or else updates her degree of belief that p . Now, according to (1), G 's warrant is a function of the warrant of its members. However, since m is a member of G , then since one of G 's members becomes better positioned epistemically towards p , G itself appears to become more warranted in believing that p . It appears, then, that G has improved its epistemic position with regards to p by simply testifying to one of its own members, and without acquiring any new evidence. In turn, G appears to become a better source of testimony with regards to p . G and m , then, have engaged in a process of groupstrapping.

Consider the following example of groupstrapping in action: an individual wants to know information about crime statistics in America. There is nothing in principle epistemically illicit about seeking out such information from a group: independent research groups that collect such

⁴ Again, depending on how we conceive of group warrant we might include other relevant factors. For instance, if group decisions are made as a result of a particular decision-making procedure, such a procedure would be a relevant factor in determining group warrant.

statistics, for instance, can be good sources of information on such matters. In this case, however, the individual appeals to testimony from a self-membership group: they log into Facebook, and find that one of the groups they are a member of – “True American Patriots for Freedom” – states that Muslim immigrants are responsible for the majority of crime in America⁵. This information is endorsed by other members of the group, in one form or another (perhaps posts within the group expressing this message have received numerous “likes”, or have been extensively shared and commented on), and dissenting views, should they exist, are shouted down. As a result, the individual forms the relevant belief, and endorses it, perhaps in the form of expressing their view within the group or adding to the perceived credibility of the relevant posts by signaling their agreement with the view. The member thus takes themselves to have acquired a new warranted belief, and with new support from one of its members, the group then appears to gain additional credibility without the acquisition of any additional evidence.

There is reason to think that groupstrapping is ubiquitous. First, it is commonplace for individuals to say that they receive their information from self-membership groups, especially when it comes to social media: according to the Pew Research Center, for example, two-thirds of Americans “get at least some of their news on social media”, with 45% using Facebook as a source of news (Shearer and Gottfried, 2017). These data do not, however, indicate whether an individual is merely using the relevant social media as platforms in which they receive information from other individuals, or if groups within those social media, or the social media themselves, are being appealed to as a source of testimony. For example, there is a difference between logging onto Facebook and receiving information from a trustworthy source through the

⁵ This example is meant to be fictional: any resemblance to groups of the same or similar name is purely coincidental. This claim is also obviously false.

website and consulting a self-membership group. There is reason to think, however, that even when it comes to information that appears to be disseminated by an individual in social media groups, that information tends to be accepted to the extent to which it is approved by the relevant community. It seems that information disseminated via “tweets” on Twitter, for example, will be more likely to be accepted the more “hearts” it receives, and posts on Facebook will be more likely to be accepted the more “likes” it receives. “Hearts” and “likes” are markers that the relevant bit of information is endorsed by the group itself, and it is this group endorsement, rather than any specific qualifications of the individual disseminating the information in the first place, that makes it more likely to be accepted. Highly endorsed views will then be taken to be the view of the group itself: one can log onto a social media platform and expect to find the expression of familiar views, and will also expect to find that dissenting views will not be highly endorsed. We can, then, take claims that someone received their information “from Facebook” or “from Twitter” in a literal sense.

The phenomenon of groupstrapping as I am presenting it here is related to a problem of belief formation and updating that has been discussed in several other debates in epistemology, namely that of evidential double-counting. At its most basic form, one double-counts some evidence E as support for a proposition P when one takes two instances of E to count as different pieces of evidential support towards P. Consider a toy example: if you believe that P only because I told you that P, you then telling me that P does not constitute any additional reason for me to believe that P. Similarly, double-counting is the phenomenon behind the familiar informal fallacy of appeals to popularity: the fact that a large number of people believe that P is not itself an indication that P is well-supported if all those individual beliefs are based on the same evidence (in such cases the same evidence may be counted more than twice).

Double-counting can occur in different ways. For example, Thomas Kelly considers the phenomenon of double-counting in arguments concerning rational peer disagreement. Here is the problem as Kelly presents it: say that we both have the same first-order evidence *E*, and are trying to decide whether to believe or disbelieve hypothesis *H* on that basis. Should I take the fact that you, a peer, believe that *H* on the basis of that evidence to be a reason to believe that *H*? Kelly thinks it should not count as such a reason. After all, since *you* would not consider the fact that *you* believe *H* to be a reason to believe it, neither should I. To do otherwise would engage in a form of double-counting: since both my belief that *H* and the fact that you believe that *H* are based on the same evidence *E*, to consider your belief that *H* to be additional evidence would be to count *E* twice.

As I mentioned earlier in the discussion of what makes epistemic bubbles and echo chambers epistemically pernicious, double-counting can also occur between members within groups. Again, this kind of double-counting might occur as a mere appeal to popularity, in which shared beliefs are based on a common set of evidence. As Nguyen (2018) notes, however, double-counting can occur within a group even when evidence is not completely shared between members of said group:

Suppose that I believe that the Paleo diet is the greatest diet of all time. I assemble a Facebook group called ‘Great Health Facts!’ and fill it only with people who already believe that Paleo is the best diet. The fact that everybody in that group agrees with me about Paleo shouldn’t increase my confidence level one bit. They’re not mere copies – they actually might have reached their conclusions independently – but their agreement can be entirely explained by my method of selection. The group’s unanimity is simply an echo of my selection criterion.

It is not obvious that Nguyen is entirely correct here: if someone really has reached their conclusions independently, or if they have new evidence or reasons that I had not considered that support my belief, it seems that I may very well be warranted in increasing my confidence on the basis of a group member's agreement. The situation that Nguyen envisions, though, is perhaps not a double-counting of first-order evidence (that is, evidence that speaks directly in favor of a proposition), but instead a double-counting of *agreement*. The problem, then, is that while someone with different evidence agrees with you may warrant an increase in confidence in your belief, the fact that a group consisting of such individuals is unanimous in their belief is not additional reason to increase confidence in your belief. To count individual and group agreement as separate pieces of evidence would be to double-count evidence in the form of double-counting agreement.

Groupstrapping, as I am proposing it here, may then be classified as a different type of double-counting, although one that is less straightforward than the types mentioned above. As I describe above, groupstrapping is a general phenomenon that exists between a member of a group and a group itself, in which a member illicitly increases her warrant in a belief by appealing to the testimony of a group, one whose epistemic position towards a proposition she is at least partially responsible for determining. When an individual member appeals to the testimony of a self-membership group, the member and the group will likely not have overlapping evidential bases: depending on how we determine what the group is warranted in believing, the evidence possessed by the member will likely be merely a subset of that evidence possessed by the group. Thus the relationship between a group and a member is not simply one of counting the same first-order evidence twice.

Groupstrapping will also not simply be a case of double-counting agreement: when the testimony of a group is not an expression of mere consensus of unanimous individuals within a group, then by appealing to a group's testimony one will not risk counting agreement twice. It is of course possible that members of groups can appeal to mere consensus in the way that Nguyen describes, and it is possible that in some cases the information expressed by a group may merely represent the expression of consensus from the members of the group. Whether type of double-counting occurs then depends on the nature of the group in question: again, to reuse Pettit's terminology, I take it that epistemically pernicious groups can come in both in the form of aggregates and integrations. The kind of group that Nguyen has in mind, for example, seems to be of the aggregate variety. In such a case one would illicitly increase one's confidence in one's belief by appealing to the aggregate, and thereby also contribute to the apparent warrant of the group itself: this variety of groupstrapping would then involve a kind of agreement double-counting in the way that Nguyen describes. But it is a special case: when group beliefs are not mere expressions of unanimity, but are instead determined by an interplay of the weighing of views expressed within the group, the mandate of that group, and perhaps a particular decision-making procedure, by appealing to the testimony of a group one is not engaged in straightforward double-counting in either the first-order or agreement sense. Furthermore, it seems that epistemic bubbles and echo chambers rarely involve the expression of truly unanimous beliefs. This is because it is likely that different members of the group have different (although perhaps overlapping) sets of evidence and reasons, that the views of members are not given equal weight, and that the views of more authoritative members will be weighed more heavily than those deemed to have less authority.

So far I have argued that while the problems created by epistemically pernicious groups are standardly thought of as a function of the effects of individuals within a group, groupstrapping represents a problem created by the effects of the group itself. Groupstrapping does not, however, merely label one more way in which one can increase confidence in one's belief to an unwarranted degree. Rather, the existence of groupstrapping can help us identify both reasons why epistemically pernicious groups form in the first place, as well as why they tend to be difficult to break out of. I turn to this next.

4. Groupstrapping and the Perceived Credibility of Epistemically Pernicious Groups

One problem concerning epistemically pernicious groups is how they form in the first place. Above we have seen that one such cause, at least as it applies to online groups, is digital curation, insofar as one tends to be exposed to similar sources and like-minded individuals as a result of algorithms that monitor one's browsing habits. However, while the fact that others agree with us does seem likely to result in an increased confidence in our beliefs, it is not clear why the fact that someone whom I know nothing about, concerning some non-trivial information, should increase my confidence to a significant degree. For example, the fact that an anonymous individual online agrees with me does not do much to make me feel more confident that our shared belief is true: this person could have no idea what they are talking about, they could hold a variety of false or reprehensible beliefs, they may be lying or generally untrustworthy, etc. Many of our interactions online occur with those about whom we know essentially nothing. So why should it be the case that agreement, even if it is widespread, should result in increased confidence to the extent that is found in epistemically pernicious groups?

Indeed, it is a problem in discussions of testimony generally what the requirements are for the acquisition of a warranted belief on the basis of testimony. Without delving too deeply into these issues here, while there certainly seem to be cases in which it is relatively unimportant what a recipient of testimony knows about the testifier – it seems plausible that one can, for example, come to know the time or directions on the basis of the testimony of strangers about whom we one knows next to nothing – we often expect that when it comes to matters that are more complex or significant that one know at least something about the testifier⁶. We should not (continuing an earlier example) expect that a randomly selected individual online should have any special knowledge about crime statistics in America, and in seeking out this information we would expect that one consult someone who is at least perceived to have some epistemic credibility in such matters.

Here, then, is the first way in which recognizing the groupstrapping problem can help us make better sense of how epistemically pernicious groups form: a group can be perceived to possess a degree of credibility that is not possessed by individual members, especially when those individuals are anonymous. While I have no particular reason to trust an opinion expressed by a pseudonym online, I may have more reason to trust the views expressed or endorsed by a group. That a group is presenting a specific view at all is perhaps better reason to believe that such view is correct, since the formation of a group implies that not only is the view widely shared, but that it is held by people who have reason to believe it. That a group expresses or endorses a view, rather than an individual, can appear to be a better reason to accept that view. When the perceived credibility of the group comes about as the result of groupstrapping,

⁶ This is so even for those who hold that we have default reason to rely on testimony: even if this is the case, this reason is defeasible, and thus we may still require that we have additional reason to trust the source of testimony when matters are important.

however, that group risks becoming epistemically pernicious, as both group and member confidence will outrun group and member warrant, respectively.

That groups obtain high degrees of perceived credibility as a result of groupstrapping can in part explain why epistemically pernicious groups are difficult to break out of, where “breaking out of” such a group can involve changing one’s mind on a central view endorsed by said group, or by getting one to take an alternate source of information seriously. Discussions of epistemically pernicious groups tend to claim that it is easier to break out of epistemic bubbles than echo chambers. As Nguyen (2018) has argued, since epistemic bubbles merely exclude views instead of actively reinforcing them, “bursting the bubble” requires that one simply expand the range of sources from which one seeks information. Similarly, Pariser (2011) argues that the solution to bursting an epistemic bubble is to add variety to one’s information-seeking routines, specifically by visiting different types of websites that one would not normally visit (223). Sunstein (2017) adds that epistemic bubbles need to be addressed on a more structural level: he proposes that “a well-functioning system of free expression” needs to facilitate both an individual’s exposure to unplanned and unexpected encounters to new information, as well as the opportunities for people to have shared experiences, so that they might be able to better understand each other on the basis of those experiences.

From what we have seen, however, it does not seem that these kinds of strategies would help when it comes to all types of epistemically pernicious groups. Again, this is due to the fact that some such groups actively discredit outside sources of information: if one is extremely confident both in one’s belief and the sources of one’s information, and if one assumes outside sources are deceptive, any kind of unplanned or unexpected encounters with conflicting information will be met not with interest in a different point of view, but with skepticism and

disdain for someone who is providing false information. This will occur even when those with marks of expertise attempt to present conflicting information to those within such a group. Consider, for example, those who refuse to believe in human-produced climate change, where even the testimony of an overwhelming majority of experts is not enough to change their minds. Climate change deniers will often appeal to the testimony of the members of one's group in order to support their beliefs. There is clearly a sense in which the practice of relying on the testimony individuals with no extensive knowledge or expertise in a relevant subject matter, as opposed to the experts, is irrational. It is easier to make sense of this behavior, however, if we think of an individual within an epistemically pernicious group as appealing not only to other individuals, but instead to the group itself. When such a group is taken to carry a mark of credibility, a member then takes themselves to have a better basis for rejecting dissenting views, even when those views are presented by experts.

It is also the perceived credibility of a group that immunizes it from criticism directed towards its individual members. That a member of a group is shown to lack authority, or to hold their beliefs without good reason, or to hold other irrational beliefs, or reprehensible beliefs, etc., will then not necessarily have any effect on the perceived credibility of the group itself, since it is the group, and not the individual, which is taken to be credible in its own right. Consider, for instance, those who hold the belief that vaccines are a cause of autism. Despite the originator of the view, Andrew Wakefield, having been widely discredited and stripped of his medical license, those within the epistemically pernicious group of "anti-vaxxers" maintain their belief. It is hard to make sense of this behavior if members of the group were merely relying on the testimony of individuals within that group, or even the once leader of said group. There is no doubt a sense in which the discrediting of Wakefield was taken by members of the anti-vaxxers group to be

illegitimate, and thus may still view Wakefield as an expert. Regardless, maintaining a belief that someone is still an authority when they have been comprehensively discredited by the scientific community is straightforwardly irrational. Again, we can, I think, make better sense of why members of this epistemically pernicious group maintain beliefs derived from sources that have been discredited: the group itself is taken to display a mark of credibility, and is treated as a source of information, rather than the discredited individual themselves.

5. Group Credibility, and Conclusion

My purpose here has been to establish that epistemically pernicious groups are pernicious not simply in virtue of the effects of the interactions of like-minded individuals, but are also partially a product of a problem with the way that members of groups rely on the groups themselves as sources of information, a process I have called groupstrapping. I argued that groupstrapping can also help explain why epistemically pernicious groups are difficult to break out of: these groups can survive attempts at discrediting views held by members since such members can appeal to the group itself for support, a group that is perceived to be more credible than it actually is.

Additionally, while we saw arguments earlier that epistemic bubbles can be burst by diversifying one's sources of information, the existence of the groupstrapping problem suggests that these bubbles may not be so easily burst after all. Again, this is because when the epistemic bubble is itself treated as a source of information it can appear to become a more credible source than other sources that offer a different viewpoint. Diversifying one's sources of information will not be effective in changing one's mind if one takes oneself to have a more credible source on their side.

While I will not attempt to present a solution to the problem of epistemically pernicious groups, I think we can, as a result of the discussion here, note that any such solution will need to address both the way that individuals rely on and influence each other, as well as how such groups are related to their members. One general way that we can approach the problem of breaking out of an epistemically pernicious group, then, could involve addressing the credibility of the group itself, not so much in terms of whether the views expressed by the group are true or false, but rather in terms of where the perceived credibility of the group is derived. One may be less likely to attribute credibility to a group, then, if one realizes that said credibility is produced through groupstrapping. While it is good advice in general to make sure that the source of one's information is actually credible, we do not, I think, note often enough how we should evaluate the credibility of groups, and especially self-membership groups. Of course, epistemic factors constitute only one part of the explanation for the formation and persistence of epistemically pernicious groups (social and political factors are undeniably salient, as well), and thus any solution to the contemporary problem of the proliferation of such groups will necessarily be multifaceted. However, as I have argued here, one of those facets ought to focus on how members of epistemically pernicious groups rely on said groups as sources of information.

Works Cited

Barbera, Pablo, John T. Jost, Jonathan Nagler, Joshua A. Tucker, and Richard Bonneau.

“Tweeting from left to right: is online political communication more than an echo chamber?”

Psychological Science 26, no. 10 (2015): 1531-1542.

Baumgaertner, Bert. “Yes, no, maybe so: a veritistic approach to echo chambers using a

trichotomous belief model.” *Synthese* 191 (2014): 2549-2569.

Jamieson, K., & Cappella, J. *Echo chamber: Rush Limbaugh and the conservative media establishment*. New York: Oxford University Press, 2009.

Justwan, Florian, Bert Baumgaertner, Juliet E. Carlisle, April K. Clark and Michael Clark.

“Social media echo chambers and satisfaction with democracy among Democrats and Republicans in the aftermath of the 2016 US elections.” *Journal of Elections, Public Opinion and Parties* (2018): 1-19.

Kelly, Thomas. “Peer Disagreement and Higher Order Evidence.” In Alvin I. Goldman & Dennis Whitcomb (eds.), *Social Epistemology: Essential Readings*. Oxford University Press, 2010: pp. 183-217.

Lackey, Jennifer. “What is Justified Group Belief?” *Philosophical Review* 125 no. 3 (2016): 341-396,

Nguyen, C. Thi. “Escape the Echo Chamber.” *Aeon*, April 9, 2018. <https://aeon.co/essays/why-its-as-hard-to-escape-an-echo-chamber-as-it-is-to-flee-a-cult>.

Pariser, Eli. *The Filter Bubble*. New York: Penguin Books, 2011.

Pettit, Philip. "Groups With Minds of their Own." In Alvin I. Goldman & Dennis Whitcomb (eds.), *Social Epistemology: Essential Readings*. Oxford University Press, 2010: pp. 167-194.

Ritchie, Katherine. "What are groups?" *Philosophical Studies* 166 (2013): 257-272.

Shearer, Elisa and Jeffrey Gottfried. "News Use Across Social Media Platforms 2017." *Pew Research Center*, September 7, 2017. <http://www.journalism.org/2017/09/07/news-use-across-social-media-platforms-2017>.

Sunstein, Cass R. *#Republic*. Princeton: Princeton University Press, 2017.

Tollefsen, Deborah. "Group testimony." *Social Epistemology* 21, no. 3 (2007):299 – 311.