

ECHO CHAMBERS AND EPISTEMIC BUBBLES

C. THI NGUYEN

cnguyen@uvu.edu

ABSTRACT

Discussion of the phenomena of post-truth and fake news often implicates the closed epistemic networks of social media. The recent conversation has, however, blurred two distinct social epistemic phenomena. An *epistemic bubble* is a social epistemic structure in which other relevant voices have been left out, perhaps accidentally. An *echo chamber* is a social epistemic structure from which other relevant voices have been actively excluded and discredited. Members of epistemic bubbles lack exposure to relevant information and arguments. Members of echo chambers, on the other hand, have been brought to systematically distrust all outside sources. In epistemic bubbles, other voices are not heard; in echo chambers, other voices are actively undermined. It is crucial to keep these phenomena distinct. First, echo chambers can explain the post-truth phenomena in a way that epistemic bubbles cannot. Second, each type of structure requires a distinct intervention. Mere exposure to evidence can shatter an epistemic bubble, but may actually reinforce an echo chamber. Finally, echo chambers are much harder to escape. Once in their grip, an agent may act with epistemic virtue, but social context will pervert those actions. Escape from an echo chamber may require a radical rebooting of one's belief system.

INTRODUCTION

Something seems to have gone wrong with the flow of information. Recent analyses of Facebook feeds and Twitter networks reveal that their user's informational input is being radically filtered, that users are being exposed largely to arguments and views with which they already agree (Saez-Trumper *et al.* 2013; An *et al.* 2014). What's more, whole segments of the population have dismissed the mainstream media as corrupt and untrustworthy. Many of us have started to wonder: are we trapped in echo chambers of our own making?¹

The recent conversation, however, has blurred two distinct, but interrelated, social epistemic phenomena, which I will call *epistemic bubbles* and *echo chambers*. Both are problematic social structures that lead their members astray. Both reinforce ideological separation. But they are different in their origins, mechanisms for operation, and avenues for treatment. Both are structures of exclusion – but epistemic bubbles exclude through

1 An earlier version of this article appeared as “Escaping the Echo Chamber” in *Aeon Magazine* (Nguyen 2018b). That earlier version was written with a general audience in mind. The present article contains more carefully articulated versions of the core definitions and arguments.

omission, while echo chambers exclude by manipulating trust and credence. However, the modern conversation often fails to distinguish between them.

Loosely, an *epistemic bubble* is a social epistemic structure in which some relevant voices have been excluded through omission. Epistemic bubbles can form with no ill intent, through ordinary processes of social selection and community formation. We seek to stay in touch with our friends, who also tend to have similar political views. But when we also use those same social networks as sources of news, then we impose on ourselves a narrowed and self-reinforcing epistemic filter, which leaves out contrary views and illegitimately inflates our epistemic self-confidence.

An *echo chamber*, on the other hand, is a social epistemic structure in which other relevant voices have been actively discredited. My analysis builds on Kathleen Hall Jamieson and Frank Capella's work, with some philosophical augmentation. According to Jamieson and Capella, an echo chamber's members share beliefs which include reasons to distrust those outside the echo chamber. Echo chambers work by systematically isolating their members from all outside epistemic sources (Jamieson and Cappella 2008: 163–236). This mechanism bears a striking resemblance to some accounts of cult indoctrination. By discrediting outsiders, echo chambers leave their members overly dependent on approved inside sources for information. In epistemic bubbles, other voices are merely not heard; in echo chambers, other voices are actively undermined. (This is a conceptual distinction; a community can practice both forms of exclusion to varying degrees.)

The contemporary discussion has been mostly focused on the phenomenon of epistemic bubbles. Cass Sunstein's famous discussions of group polarization, extremism, and the Internet largely focus on matters of constricted information flow and omitted viewpoints (Sunstein 2001, 2009a, 2009b). Eli Pariser's *The Filter Bubble* focuses entirely on filtration effects from personalization technology, as in Google searches, and self-selected informational networks, as in Facebook (Pariser 2011). Popular conversation has tended to follow Pariser's focus on technologically mediated filtration. The term "echo chamber" has, in recent usage, been reduced to a synonym for such bubbles and their constricted information flow. When the specifically trust-oriented manipulations of echo chambers are discussed, they are usually lumped in with epistemic bubbles as part of one unified phenomenon.

This is a mistake; it is vital to distinguish between these two phenomena. An epistemic bubble is an epistemic structure emerging from the informational architecture of communities, social networks, media, and other sources of information and argument. It is an impaired informational topology – a structure with poor connectivity. An echo chamber, on the other hand, is an epistemic structure created through the manipulation of trust; it can exist within a healthy informational topology by adding a superstructure of discredit and authority. I hope to show, contra the recent focus on epistemic bubbles, that echo chambers pose a significant and distinctive threat – perhaps even a more dangerous one – that requires a very different mode of repair. Furthermore, echo chambers can explain what epistemic bubbles cannot: the apparent resistance to clear evidence found in some groups, such as climate change deniers and anti-vaccination groups.

It may be tempting to treat members of echo chambers as mere sheep, and accuse them of problematic acquiescence to epistemic authority. But that accusation relies on an unreasonable expectation for radical epistemic autonomy. Contemporary epistemology, especially social epistemology, has taught us that trust in others is necessary and ineradicable. We are, as John Hardwig says, irredeemably epistemically dependent on

each other (Hardwig 1985, 1991). Echo chambers prey on our epistemic interdependence. Thus, in some circumstances, echo chamber members do not have full epistemic responsibility for their beliefs. Once one is trapped in an echo chamber, one might follow good epistemic practices and still be led further astray. And some people can be trapped in echo chambers because of circumstances beyond their control – for example, they can be raised in them. Which leads to the most important questions: How can one tell if one is in an echo chamber? And how can one escape? I will argue that, for those trapped within an echo chamber, prospects for detection are poor and the escape path daunting. Detecting and escaping from echo chambers will require a radical restructuring of a member’s relationship to their epistemic past, which may be more than we can reasonably expect of one another.

I. EPISTEMIC BUBBLES

Let’s start with a simple picture of how many of us conduct our epistemic lives right now. I get much of my news via Facebook. I have selected most of my Facebook friends for social reasons; they are my friends and colleagues. A significant conduit for my learning about events in the world is through people re-posting articles that they have found newsworthy or interesting. When I go outside of Facebook, I usually turn to sources which, by and large, are affiliated with my own political beliefs and intellectual culture.

This process imposes a filter on my uptake of information. Selection and exclusion are not bad in and of themselves – the world is overstuffed with supposed sources of information, many of them terrible. The better my filter, the more it will focus my attention on relevant, useful, and reliable information. But the bare fact that each individual member of the system is themselves reliable will not guarantee any broadness or completeness of coverage. Suppose, for example, that my social network was composed entirely of intelligent, reliable professors of aesthetics whose interests were largely focused on new developments in opera, ballet, and avant-garde contemporary art. Through this system, I might learn about all the exciting new developments in the New York Art scene, but entirely miss, say, relevant developments in rap, or the fact that my country was slowly sliding into fascism. The system lacks what Goldberg calls coverage-reliability – the completeness of relevant testimony from across one’s whole epistemic community (Goldberg 2011: 93–4). Bad coverage can not only leave out relevant facts and evidence; it can also fail to bring relevant arguments to our attention. Thus, bad coverage can starve us of adequate exposure to relevant arguments. Notice that bad coverage is an epistemic flaw of epistemic systems and networks, not of individuals.

I can now specify my use of “epistemic bubble” with greater precision. An epistemic bubble is *a social epistemic structure which has inadequate coverage through a process of exclusion by omission*. Epistemic bubbles form by leaving out relevant epistemic sources, rather than actively discrediting them. There are two primary forces encouraging this omission. First, there is an epistemic agent’s own tendency to seek like-minded sources. This phenomenon is sometimes called “selective exposure” by social scientists (Nelson and Webster 2017). In many contemporary cases, such as Facebook, the process of omission can occur inadvertently. I typically put people in my Facebook feed for social reasons – because I like them or I find them funny. But social selection does not guarantee good coverage reliability; in fact, the typical bases of social selection are inimical to good

coverage reliability.² We usually like people who are similar to us, and such similarity makes coverage gaps more likely. Friends make for good parties, but poor information networks. We now have a straightforward account of one way in which epistemic bubbles can form. We can build a structure for one set of purposes – maintaining social relations, finding – and then proceed to use it for an entirely different purpose, for which it functions badly: information gathering.

Second, there are the processes by which an epistemic agent’s informational landscape is modified by other agents. This might include, say, systematic censorship or media control by the state or other actors. The most worrisome of these external forces, at the moment, seems to be the *algorithmic personal filtering* of online experiences (Pariser 2011; Watson 2015). Internet search engines, for example, will track personal information for each particular user, adapting their search results to suit each user’s interest. Certainly, newspapers and other traditional media technologies do place external filters on their readers, but the modern instantiation is particularly powerful and troubling. As Boaz Miller and Isaac Record argue, Internet technologies create hyper-individualized, secret filters. The secrecy is particularly threatening. Many users do not know about the existence of algorithmic personal filtering. Even amongst those that do, most do not have access to the particularities of the algorithms doing the filtering; thus, the very opacity of the process makes it harder for a user to successfully evaluate and epistemically compensate for such filtering (Miller and Record 2013). Thus, most users significantly underestimate the degree to which their exposure to information, via search results, has already been tailored to present information to which the user will already be amenable.

Both the agent-driven process of selective exposure, and the externalities of algorithmic filtering, contribute to the creation of epistemic bubbles. I introduce the term “epistemic bubble” here to indicate a broader set of phenomena. Pariser introduced the term “filter bubble” to refer specifically to technologically mediated filtering, especially via algorithmic matching. Epistemic bubbles are those structures which omit relevant voices by any means, technological or otherwise. Epistemic bubbles include filter bubbles, but also non-technological selection processes, such as physically sorting oneself into neighborhoods of like-minded people (Bishop 2009).

The account I’ve given of epistemic bubbles focuses on the way they omit relevant information, but that omission can also threaten us with *bootstrapped corroboration*. Users of social networks and personalized search technologies will encounter agreement more frequently and so be tempted to over-inflate their epistemic self-confidence. This danger threatens because, in general, corroboration is often a very good reason to increase one’s confidence in the relevant beliefs (Nguyen 2010, 2018a). But corroboration ought to only have weight if it adds something epistemically, rather than being a mere copy. To borrow an example from Wittgenstein: imagine looking through a stack of identical newspapers and treating each next newspaper headline saying *p* as a reason to increase your belief that *p* (Wittgenstein 2010: 100). This is clearly a mistake; the fact that a newspaper claims that *p* has some epistemic weight, but the number of copies of that newspaper one encounters ought not add any extra weight. Similarly, imagine speaking to a bunch of acolytes of Guru Jane, who repeat anything that Guru Jane says without any further thought. The fact that all these acolytes repeat Guru Jane’s testimony should add no extra weight. So long as the

2 For an overview of empirical research on personal similarity and polarization, see Sunstein (2009a: 83–5).

disciplines repeat anything Guru Jane says – so long as they are mere conduits for information, rather than sources of information – they are simply another sort of copy.

But copying isn't the only route to a problematic form of non-independence. Suppose I believe that the Paleo diet is the best diet. I proceed to assemble a body of peers who I trust precisely because they also believe that Paleo is the best diet. In that case, the existence of perfect agreement on Paleo's amazingness throughout that group ought to count for far less than it might for other groups that I had not assembled on that basis. Even if all the group members arrived at their beliefs independently, their agreement is already guaranteed by my selection principle. To the degree that I have pre-selected the members in my epistemic network based on agreement with some set of beliefs of mine, then their agreement with that set of beliefs and any other beliefs that it entails ought to be epistemically discounted.³ If we fail to so discount, we are ignoring a pernicious hidden circularity in our corroborative process. It is easy to forget to discount because the bootstrap here is obscured by time and interface. But we must actively adjust for the increased likelihood of agreement inside our bubbles, or we will unwarrantedly bootstrap our confidence levels.⁴

To summarize: an epistemic bubble is an epistemic network that has inadequate coverage through a process of exclusion by omission. That omission need not be malicious or even intentional, but members of that community will not receive all the relevant evidence, nor be exposed to a balanced set of arguments.

2. ECHO CHAMBERS

Luckily for us, epistemic bubbles are relatively fragile. Relevant sources have simply been left out; they have not been discredited. It is possible to pop an epistemic bubble by exposing a member to relevant information or arguments that they have missed. Echo chambers, on the other hand, are significantly more robust.

My analysis here combines empirical work and analysis from Jamieson and Cappella on the nature of right-wing echo-chambers with recent insights from social epistemology. Jamieson and Cappella studied echo chambers built around particular charismatic personalities – Rush Limbaugh and the news team of Fox News, and certain other members of conservative talk radio. Their data and analysis suggest that Limbaugh uses various methods to actively isolate his community of followers from other epistemic sources. Limbaugh's consistent attacks on the "mainstream media" serve to discredit all potential

3 This is a relative of the problem, from statistics, of failing to identify dependent variables.

4 I am not claiming, as Alvin Goldman does, that all cases of non-independent testimony should be discounted (Goldman 2001: 98–104). Goldman's analysis, and the principle of non-independence, have been significantly challenged (Coady 2006; Lackey 2013). David Coady and Jennifer Lackey have demonstrated that we can trust the weight of non-independent agreement when we have good reason to think that the non-independent agreeers had good epistemic meta-reasons for agreeing. For example, suppose that all scientists agreed that climate change was coming. Their agreement is non-independent – the majority of these scientists have not analysed the data for themselves, but trust the expert specialists in climate change. But still, the weight of numbers matters here because the trusting scientists have good epistemic reasons for picking who to trust. But in the case I've described, the weight of numbers here does not emanate from a good epistemic process. In other words, my claim doesn't depend on the claim that all forms of non-independence are problematic, only that some are, for the reasons I adduce.

sources of knowledge or testimony besides Limbaugh and a select inner-cadre of other approved sources (Jamieson and Cappella 2008: 140–76). Limbaugh also develops what they call a private language, full of alternate meanings for familiar terms and new jargon (for example, “SJWs”), in order to exaggerate the insularity and separateness of the in-group. Finally, Limbaugh provides counter-explanations for all contrary views, intended not only to attack each particular view, but also to undermine the general trustworthiness and integrity of anybody expressing a contrary view. The resulting world-view is one of highly opposed forces; once one has subscribed to Limbaugh’s view, one has reason to think that anybody who does not also subscribe is actively opposed to the side of right, and thereby morally unsound and so generally untrustworthy (Jamieson and Cappella 2008: 177–90). Jamieson and Cappella suggest that this makes a follower dependent on a single source or group of sources, and makes them highly resistant to any outside sources. They offer the following definition of an echo chamber: an echo chamber is a bounded and enclosed group that magnifies the internal voices and insulates them from rebuttal (Jamieson and Cappella 2008: 76).

I will use the term “echo chamber” here following their analysis, but I adapt the definition slightly for philosophical use. 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*. Finally, echo chambers are such that *general agreement with some core set of beliefs is a prerequisite for membership*, where those core beliefs *include beliefs that support that disparity in trust*.

By “epistemic discrediting”, I mean that non-members are not simply omitted or not heard, but are actively assigned some epistemic demerit, such as unreliability, epistemic maliciousness, or dishonesty. By “amplifying epistemic credentials”, I mean that members are assigned very high levels of trust. Of course, these two processes can feed back into one another. So long as an echo chamber’s trusted insiders continue to claim that outsiders are untrustworthy, then the inner trust will reinforce the outward distrust. And so long as outsiders are largely distrusted, then the insiders will be insulated from various forms of counter-evidence and rebuff, thus increasing their relative credence. Once a sufficient disparity in credence between insiders and outsiders has been established, so long as trusted insiders continue to hold and espouse epistemically dismissive beliefs towards outsiders, then the echo chambers’ beliefs system will be extremely difficult to dislodge.

Compare this process of credence manipulation to the process of omission found in epistemic bubbles. In one standard scenario, I add others as trusted members of my epistemic network based on agreement. I am then less likely to encounter an outside voice – but when I do actually have such an encounter with an outsider, I have no background reason to dismiss them. Bubbles restrict access to outsiders, but don’t necessarily change their credibility. Echo chambers, on the other hand, work by offering a pre-emptive discredit towards any outside sources.⁵

5 Note that the kind of echo chambers here are different from those I described in Nguyen (2018a). There, I explored what I called “personal” echo chambers, which can be generated through entirely good-intentioned individual action under very special epistemic conditions, which I called a “cognitive island”. A cognitive island is any cognitive domain in which there is no possible empirical method to check whether a purported expert is really an expert. For example, the moral domain and the aesthetic

The result is a rather striking parallel to the techniques of isolation typically practiced in cult indoctrination. The standard techniques of cult indoctrination, by a traditional account, are the aggressive emotional isolation of cult members from all non-cult members, which amplifies indoctrinated member's dependency on the cult (Singer 1979; Lifton 1991; Langone 1994).⁶ New cult members are brought to distrust all non-cult members, which provides an epistemic buffer against any attempts to extract the indoctrinated person from the cult. This is nothing like how epistemic bubbles work. Epistemic bubbles merely leave their members ignorant, but ignorance can be fixed with simple exposure. The function of an echo chamber, on the other hand, is to *credentially isolate* its members through a manipulation of trust. By this, I mean that members are not just cut off, but are actively alienated from any of the usual sources of contrary argument, consideration, or evidence. Members have been prepared to discredit and distrust any outside sources; thus, mere exposure to relevant outside information will have no effect.

In fact, echo chambers can avail themselves of another epistemic protective mechanism: they can contain what I'll call a *disagreement-reinforcement mechanism*. Members can be brought to hold a set of beliefs such that the existence and expression of contrary beliefs reinforces the original set of beliefs and the discrediting story. A toy example: suppose I am a cult leader, and I have taught my followers to believe that every human except the members of our group has been infested and mind-controlled by alien ghosts from Mars. I also teach my followers that these alien ghosts from Mars hate our group for knowing the truth, and so will constantly seek to undermine our knowledge of their existence through mechanisms like calling us a 'cult' and calling us lunatics. Endre Begby has offered a careful analysis of this particular sort of disagreement-reinforcement mechanism, which he calls "evidential pre-emption". Suppose that I tell my followers to expect outsiders to falsely claim that there are no ghosts from Mars. When my followers do confront such contrary claims from outsiders, those contrary claims are exactly what they expected to hear. Thus, new contrary testimony is neutralized, because it was predicted by past beliefs. This, says Begby, functions as a kind of epistemic inoculation. There is also a secondary effect. When my followers hear exactly what I predicted, then my claims have been verified, and so my followers will have some reason to increase their trust in me. Thus, the echo chamber's belief system not only neutralizes the epistemic impact of exposure to outsiders with contrary beliefs; the existence of those contrary beliefs will actively corroborate the pre-emptor and so increase the credence level of the entire echo chamber (Begby 2017). This creates a feedback mechanism within the echo chambers. By making undermining predictions about contrary testimony, inside authorities not only discredit that contrary testimony, but increase their trustworthiness for future predictions.

domain might be cognitive islands. What I've discussed in this paper is something different – what I've called "social" echo chambers. They are malicious and essentially social structures which can exist in any empirical domain, not just on a cognitive island. I will offer a synthesis of these two conceptions of echo chambers, with an eye towards elucidating the similarities in their underlying structures, in a future work.

- 6 Note that this view of cult indoctrination is standard among mental health professionals and social workers, but has been resisted by some theorists, especially from the fields of religious studies and sociology, as a method of de-legitimizing and repressing minority religions (Robbins and Anthony 1982). For a good overview of the debate, see Coates (2010). For the purposes of my argument, it does not matter if cults actually work this way; what is interesting is that echo chambers work by the mechanism that is, under a traditional account, attributed to cults.

Once such a system of beliefs is set up, it can be very difficult to dislodge – it is self-reinforcing, bounded, and built to discount any contrary input. In fact, though my definition of echo chambers is conceptually separable from such a disagreement-reinforcement mechanism, every plausible real-world candidate for an echo chamber I’ve ever encountered includes some version of a disagreement-reinforcement mechanism. For a depressing real-world example, consider Pizzagate. Pizzagate is a conspiracy theory that boiled out of a right-wing online forum on Reddit, which included beliefs that Comet Ping Pong, a pizza restaurant, was the site of a child sex trafficking ring owned by a liberal conspiracy involving Hillary Clinton and Barack Obama. Eventually, Edgar Welch, a member of that forum, forcibly entered the pizza parlor armed with an assault rifle to investigate; when he satisfied himself that the restaurant contained no child slaves, he gave himself up to the police. The online forum, however, did not take this as contrary evidence. Instead, they leaned on their belief that the liberal conspiracy had total control of the mainstream media, and was willing to stage fake events to discredit the right-wing. The forum took Welch’s claims that there was no sex trafficking ring as evidence that Welch was a paid actor, and thus as further confirmation of the existence of a powerful cabal of liberal child sex traffickers (Mengus 2016; Vogt and Goldman 2016).

Conspiracy theories function here in a fascinating inversion to corroborative bootstrapping. In corroborative bootstrapping, the mistake is to treat problematically dependently selected insiders as if they were independent, and thus overweight their testimony. When an echo chamber uses a conspiracy theory in this manner, they are attributing a problematic form of non-independence to outsiders who are actually independent, and thereby underweighting outside testimony. An echo chamber here works by discrediting the apparent independence of, say, different climate change scientists by claiming that all their various testimonies are problematically derived from a single source. Incidentally, I am not claiming here that conspiracy theories are always or necessarily incorrect or the product of epistemic vice. As others have argued, believing in conspiracy theories isn’t bad per se, because some conspiracy theories are true (Coady 2012: 110–37; Dentith 2015, 2017). But the fact that conspiracy theories can function to reinforce the boundaries of echo chambers – though they do not necessarily do so – might explain part of conspiracy theories’ bad rap. Because of their effectiveness in setting up disagreement-reinforcement mechanisms, conspiracy theories are often conscripted as a powerful tool in the bad epistemic behavior of certain groups.

It is important to note that the epistemic mechanisms by which echo chambers work, though problematic, are not *sui generis*. They are perversions of natural, useful, and necessary attitudes of individual and institutional trust. The problem isn’t that we trust and distrust groups and institutions. In fact, we must do so. Eljiah Millgram calls it the problem of hyper-specialization. Human knowledge has splintered into a vast set of specialized fields that depend on each other. No one human can manage that information; we are forced to trust each other (Millgram 2015: 27–44).⁷ None of us is in a position to reliably

7 Though this paper relies on insights from modern work in the epistemology of testimony, I have tried to rely only on uncontroversial claims from that field, and not on the technical details of any particular view. In particular, I have attempted to construct my analysis so as to be independent of the debate on whether or not testimony is a basic source of knowledge. I have also attempted to make the paper compatible with the major accounts of trust.

identify an expert in most specialist fields outside of our own. I am, on my own, helpless to evaluate the virtues of antibiotics or the expertise of a particular doctor or surgeon. I am, instead, reliant on a vast network of institutional licensing practices in order to choose my health care sources – including journal peer review, medical board exams, university hiring practices, and the like (Nguyen 2018a). Often, I trust via what Philip Kitcher calls indirect calibration – I trust mechanical engineers because they make things that work, but I know that mechanical engineers trust applied physicists, and I know that applied physicists trust theoretical physicists, so I acquire trust through a long chain of field-wide links (Kitcher 1993: 320–3). I even use litmus tests: the fact that any person or group is in favor of, say, sexual orientation conversion therapy is enough for me to discredit them on any social or moral topics. We must resort to such tactics in order to navigate the hyper-specialized world (Nguyen 2018c).

Echo chambers function parasitically on these practices by applying discredits without regard for the actual epistemic worth of the discredited institutions or individuals. The discredit is instead applied strategically and defensively, towards all outsiders solely on the basis of their being outsiders. Once the discrediting beliefs are in place, the ensuing beliefs and actions of the echo chambers' members are surprisingly close to rational. In fact, we can easily imagine alternative scenarios in which a very similar set of beliefs were appropriate and veristic. If I were an anti-Nazi in Germany during the rise of the Nazi party, I would do well to maintain the beliefs that most people were corrupt, untrustworthy, and out to maliciously undermine my own true beliefs. But if such beliefs become implanted in an inappropriate context, they can lead their believers entirely astray.

Epistemic bubbles can easily form accidentally. But the most plausible explanation for the particular features of echo chambers is something more malicious. Echo chambers are excellent tools to maintain, reinforce, and expand power through epistemic control. Thus, it is likely (though not necessary) that echo chambers are set up intentionally, or at least maintained, for this functionality. My account thus bears some resemblance to some work on testimonial injustice and the epistemology of ignorance, but it is importantly distinct. Miranda Fricker has argued for a kind of testimonial injustice, based on a gap between actual reliability and perceived credibility. For example, says Fricker, being white and being male are both bonuses to credibility. Since credibility is a source of power, anybody with credibility will seek to increase it, using that very credibility. Thus, says Fricker, credibility gaps can be turned into epistemic tools of social oppression (Fricker 2011). Similarly, Charles Mills argues that there is an active practice of ignorance among members of oppressive groups, such as white Americans. It is to the benefit of those in power to actively ignore many aspects of the existence of oppressed groups (Alcoff 2007: 47–57; Mills 2007). My account is compatible with, but independent from, Fricker's and Mills' accounts. Echo chambers can and surely are used to maintain social oppression through enhancing credibility gaps and supporting practices of active ignorance. The systematic mistrust of an echo chambers is a powerful tool for perpetuating epistemic injustice and active ignorance. However, the concept of an echo chamber does not require that they be deployed only in political contexts, nor does it require that they be deployed only in the service of oppression. Echo chambers could potentially exist among the oppressed, and surely exist in apolitical contexts. I believe I have witnessed echo chambers forming around topics such as anti-vaccination, multi-level marketing programs, particular diets, exercise programs, liberal

activism, therapeutic methodologies, philosophies of child-rearing, particular academic sub-disciplines, and Crossfit (Weathers 2014).

3. “POST-TRUTH” AND THE POWER OF ECHO CHAMBERS

It has often been claimed, during and after the American political season of 2016, that we have entered a “post-truth era”. Not only do some political figures seem to speak with a blatant disregard for the facts, their supporters seem unswayed by reason or contrary evidence. To many, it seems as if a vast swath of the electorate has become entirely unmoored from any interest in facts or evidence. Call this the “total irrationality” explanation of the post-truth phenomenon.

But echo chambers offer an alternative explanation for the apparent post-truth mood. It seems likely that there is at least one vast partisan echo-chamber present in the political landscape. Jamieson and Cappella’s study is a decade old, but sources like Breitbart and Alex Jones’ Infowars seem like clear extensions of the same right-wing echo chamber. (Other echo chambers surely exist elsewhere on the political spectrum, though, to my mind, the left-wing echo chambers have been unable to exert a similar level of political force.) In that case, the account of echo chambers I’ve offered has significant explanatory force. The apparent “post-truth” attitude can be explained, at least in part, as the result of credence manipulations wrought by echo chambers. In healthy epistemic communities, there is something like an upper ceiling on the credence level accorded to any individual. A healthy epistemic network will supply a steady stream of contrary evidence and counterarguments; thus, no single individual or group will ever go unchallenged. Epistemic bubbles make the discovery of mistakes significantly less likely, and so tend to exaggerate the credence levels of epistemic sources inside the bubble. But when an echo chamber is in place and all outside sources have been effectively discredited, that ceiling disappears categorically. Echo chambers can create *runaway credence levels* for approved individuals. By removing disconfirmations and disconfirmation from the system through the systematic discrediting of outsiders, echo chambers can create exceptionally high – one is tempted to say unnaturally high – levels of trust. That potential for runaway credence is built right into the foundations of any echo chamber, and arises from an interaction between the two main components of any echo chamber. First, an echo chamber involves a significant disparity of trust between the insiders and the outsiders. Second, an echo chamber involves beliefs, espoused by the insiders, reinforcing that disparity. The essential features of echo chambers seem designed to self-reinforce their peculiar arrangement of trust.

Notice that epistemic bubbles alone cannot explain the post-truth phenomenon. Since epistemic bubbles work only via coverage gaps, they offer little in the way of explanation for why an individual would reject clear evidence when they actually do encounter it. Coverage gaps cannot explain how somebody could, say, continue to deny the existence of climate change when actually confronted with the overwhelming evidence. One would be tempted, then, to accuse climate change deniers of some kind of brute error. But echo chambers offer an explanation of the phenomenon without resorting to attributions of brute irrationality. Climate change deniers have entered an epistemic structure whereby all outside sources of evidence have been thoroughly discredited. Entering that epistemic structure might itself involve various epistemic mistakes and

vices – but here the story can be one of the slow accumulation of minor mistakes, which gradually embed the believer in a self-reinforcing, internally coherent, but ultimately misleading epistemic structure.

Similarly, some have suggested that, in the post-truth era, many people’s interest in the truth has evaporated. Once again, this account of echo chambers suggests a less damning and more modest explanation. An echo chamber doesn’t erode a member’s interest in the truth; it merely manipulates their credence levels such that radically different sources and institutions will be considered proper sources of evidence. This phenomenon stands in stark contrast to accounts of obfuscatory speech. Take, for instance, Orwellian double speak – deliberately ambiguous, euphemism-filled language, designed to hide the intent of the speaker (Orwell 1968). Double speak is a practice that evinces no interest in coherence, clarity, or truth. But by my account, we should expect discourse within echo chambers to be entirely different – we should expect such discourse to be crisp and clear, and to present unambiguous claims about what is the case, what secret conspiracies are in place, and which sources are to be entirely distrusted. And this is precisely what we find (Jamieson and Cappella 2008: 3–41, 140–76). Consider, for example, Breitbart’s attacks on other media sources. One article begins: “Mainstream media outlets continue to print false and defamatory descriptions of Breitbart News in a nakedly political effort to marginalize a growing competitor” (Pollak 2017). This is not the double-speak of administrators and bureaucrats – this is a clear, strident, and unambiguously worded discredit.

One might be tempted to say: but just give them the real evidence! You can’t discredit neutral evidence! But this response radically underestimates the degree of trust and social processing involved in most presentations of evidence. Except for empirical evidence I myself have gathered, all other presentations of evidence rely on trust. My belief in the reality of climate change depends on enormous amounts of institutional trust. I have not gathered the climate change evidence myself; I mostly just trust science journalists who, in turn, trust institutional credentialing systems. Even if I had been on, say, a core sampling expedition to the Arctic, I would be unable to process that information for myself, or even vet whether somebody else has properly processed it. Even the climatologist who actually processes that information must also depend on trusting a vast array of other experts, including statisticians, chemists, and the programmers of their data analysis software. Most so-called “neutral evidence” depends on long and robust chains of trust (Millgram 2015: 27–44). Members of an echo chamber have acquired beliefs which break the usual arrangements of trust.

But despite their evident explanatory force, echo chambers have been largely neglected by recent empirical research. Much of the recent research on causes of belief polarization focuses on the causal role of individual psychology, such as the tendency towards laziness in the scrutiny of one’s own beliefs (Trouche *et al.* 2016). Similarly, recent studies on climate change denial focus on studying the relationship between an individual’s stated political beliefs and their reactions to climate change information, without inquiring into the social epistemic structures in which the individuals are embedded (Corner *et al.* 2012). Famously, Dan Kahan and Donald Braman argue for the *cultural cognition* thesis – that is, that cultural commitments are prior to factual beliefs, and that non-evidentially formed cultural values inform which future presentations of evidence will be admitted as weighty (Kahan and Braman 2006). Though the values may originally come from an individual’s culture, Kahan and Braman focus

their analysis on how those acquired values function in individual reasoning to create polarization. They pay little attention to the continuing role of the contingent social structures in which the individual is embedded.

The direct literature on echo chambers and epistemic bubbles is new and relatively small, compared to the sizable literature on individual belief polarization. Unfortunately, even in that literature, echo chambers and epistemic bubbles have often been confused. They are usually addressed in the popular media together, and the terms “epistemic bubble” and “echo chamber” are typically used interchangeably (El-Bermawy 2016). The same blurring has occurred in the treatment of the phenomena in academic epistemology in the surprisingly small literature on echo chambers. For example, Bert Baumgaertner, in his analysis of echo chambers via computer modeling, lumps together under the heading “echo chamber” both Jamieson and Cappella’s account of right wing echo chambers and Eric Gilbert *et al.*’s treatment of blogs as echo chambers (Baumgaertner 2014). But where Jamieson and Cappella’s subject of study is echo chambers, Gilbert’s discussion concerns how blog communities have merely excluded relevant voices through social practices of connecting with like-minded individuals – clearly what I’m calling epistemic bubbles (Gilbert *et al.* 2009). In both popular and academic cases, further analysis has focused on epistemic bubbles; the structures of discredit involved in Jamieson and Cappella’s echo chambers have been largely neglected.⁸

For a particularly prominent example, consider Sunstein’s influential account of the relationship between Internet technologies, group polarization and the rise of political extremism. Though he professes to cover both filter bubbles and echo chambers, his work focuses almost entirely on epistemic bubble effects: constricted information flow, lack of exposure to alternate arguments, and bootstrapped corroboration (Sunstein 2009a: 1–98; 2009b: xi, 19–96). The point here is about more just than his choice of words: his subjects of analysis include, among other things, Facebook friend groups, hate groups, extremist online political forums, conspiracy theorists, and terrorist groups (Sunstein 2009a: 99–125; 2009b: 46–96). Clearly, this list includes prime candidates for both epistemic bubbles and echo chambers. But his analysis focuses almost entirely on the effects of bootstrapped corroboration and lack of exposure. For Sunstein, the primary mechanism driving polarization and extremism is the loss of truly public forums, because technology has over-empowered people’s tendency to self-select sources offering familiar views. Thus, his solution is to re-create, in the new media environment, the kind of general public forums where people might be more likely to serendipitously encounter contrary views and arguments. His solutions include government-funded public news websites with diverse coverage and voluntary work by corporations and individuals to burst their bubbles. His recommendations for repair largely have to do with increasing exposure (Sunstein 2009a: 135–48; 2009b: 19–45, 190–211). But, again, if what’s going on is actually an echo chamber effect, exposure is useless or worse.

The blurring of the two concepts has also led to some problematic dismissals of the whole cluster of phenomena. A number of recent articles in social science, communications, and media studies have argued that the worries about bubbles and echo chambers

8 One exception is the careful treatment of aesthetic echo chambers in Robson (2014), which brought Jamieson and Cappella’s work to my attention, and was instrumental in this paper’s development.

has been wildly overstated. These articles share the same argumentative pattern. First, they use the terms “filter bubble” and “echo chamber” interchangeably, and address themselves to the same cluster of phenomena as Sunstein, treating them as singular. In fact, James Nelson and James Webster conflate Jamieson and Cappella’s analysis of echo chambers and Pariser’s analysis of filter bubbles, and erroneously attribute to Jamieson and Cappella the view that political partisans only seek out and encounter media from sources with matching political alignments. That is, Nelson and Webster attribute an epistemic bubbles account to Jamieson and Cappella, where Jamieson and Cappella’s actual text is clearly an echo chambers account (Nelson and Webster 2017: 2). More importantly, these recent articles proceed to argue against the existence of filter bubbles and echo chambers by demonstrating that, through the analysis of empirical data about media consumption, most people in fact expose themselves to media from across the political spectrum. Nelson and Webster, for example, argue against Jamieson and Capella, claiming that filter bubbles and echo chambers don’t exist. Nelson and Webster support their claim with data showing that both liberals and conservatives visit the same media sites and spend comparable amounts of time at those sites (Nelson and Webster 2017: 6–7). Again, this misses the mark – this is evidence only against the existence of epistemic bubbles, and not against the existence of echo chambers. Similarly, Seth Flaxman *et al.* seek to problematize the existence of filter bubbles and echo chambers with data that social media platforms seem to actually increase people’s exposure to media from across the political divide (Flaxman *et al.* 2016). Again, these data only concern exposure and omission, and weigh only against the existence of epistemic bubbles. They say nothing about whether echo chambers exist. Echo chambers, recall, are structures of strategic discrediting, rather than bad informational connectivity. Echo chambers can exist even when information flows well. In fact, echo chambers should hope that their members are exposed to media from the outside; if the right disagreement reinforcement mechanisms are in place, that exposure will only reinforce the echo chambers’ members’ allegiance. We ought not conclude then, from data that epistemic bubbles do not exist, that echo chambers also do not exist.

We can see how crucial it is to keep these two categories distinct. Epistemic bubbles are rather ramshackle – they go up easily, but they are easy to take down. Since there is no systematic discrediting of outsiders, simple exposure to excluded voices can relieve the problem. Echo chambers, on the other hand, are much harder to escape. Echo chambers can start to seem almost like living things – the belief systems provide structural integrity and resilience. Mere exposure will be ineffective. Jamieson and Cappella offer evidence of this effect: once listeners are caught in Rush Limbaugh’s language, framing, and discrediting of the mainstream media, their beliefs can survive frequent contact with contrary viewpoints. Limbaugh’s technique, say Jamieson and Cappella, serves to insulate and inoculate his audience from being affected by exposure to contrary viewpoints (Jamieson and Cappella 2008: 163–90). In fact, if the appropriate disagreement-reinforcement mechanisms are in place, exposure will simply strengthen the attacked belief systems. Thus, an outsider’s attempt to break an echo chamber as if it were a mere bubble is likely to backfire and reinforce the echo chamber’s grip.⁹

9 Sunstein does briefly note the empirical data for the disagreement-reinforcement effect in passing, but then seems to ignore it in proposing his solutions (Sunstein 2009a: 54–5).

4. RESPONSIBILITY AND THE POSSIBILITY OF ESCAPE

So what, then, are the epistemic responsibilities of an agent to discover whether they are in one of these social epistemic traps, and what are their prospects for actually discovering their predicament and successfully escaping? To answer this, we must consider two distinct questions:

The escape route question: Is there any way out of an echo chamber or epistemic bubble?

The escape responsibility question: Could one behave epistemically virtuously, and yet still remain caught within an echo chamber or epistemic bubble? In other words, to what degree is an epistemic agent embedded within such a structure blameworthy, or blameless, for the faultiness of their beliefs?

The first question asks about the possible existence of an escape route. The second asks whether there is an escape route that we might reasonably expect an epistemically virtuous agent to discover and enact. These are distinct questions, because an escape route might turn out to be possible, but so difficult to discover or use that it was beyond what we might reasonably expect of an agent of moderate epistemic virtue.

For epistemic bubbles, the answers are straightforward. As I've argued, epistemic bubbles are quite easy to shatter. One just needs exposure to excluded information. Insofar as that information is available, but simply not part of one's standard network, then members of epistemic bubbles are failing to live up to their epistemic duties, which include proactively gathering relevant data. To translate into contemporary terms: if you're subject to an epistemic bubble because you get all your news from Facebook and don't bother to look at other sources, then you are, indeed, blameworthy for that failure. If one finds the language of epistemic virtues and vices appealing, then we can say that members of epistemic bubbles are committing the vice of epistemic laziness.

Answering these two questions is much more difficult for echo chambers. Recall: where encountering excluded voices and evidence will shatter an epistemic bubble, such encounters are likely to reinforce an echo chamber. Let's grant that intentionally constructing an echo chamber, as Jamieson and Cappella claim that Rush Limbaugh did, is epistemically (and morally) blameworthy. Furthermore, actively entering an echo chamber seems epistemically blameworthy in many circumstances. For agent in full possession of a wide range of informational sources, to abandon most of them and place their trust in an echo chamber for, say, an increased sense of comfort and security, is surely some form of epistemic vice. There is some evidence that this may be the case; Jamieson and Cappella suggest that people enter echo chambers for the sake of the community bonds and the sense of belonging to an in-group (Jamieson and Cappella 2008: 180).

But there are many cases in which the agent seems plausibly blameless. Imagine a person raised in an echo chamber. Their earliest epistemic contacts – let's say their parents, relatives, and close family friends – are all firmly committed members of the echo chamber. Suppose that the child is either home-schooled by those echo chamber members or sent to a school that reinforces the beliefs of that particular echo chamber. I take it that it is reasonable for a child to trust their parents and those of seeming epistemic authority, and that a child is epistemically blameless for having done so (Goldberg 2013). Thus, when that child eventually comes into contact with the larger epistemic world – say, as

a teenager – the echo chamber’s beliefs are fully in place, such that the teenager discredits all sources outside of their echo chamber.

It seems, at first glance, that our teenager could be acting very much like a reasonable epistemic agent. They could, in fact, be epistemically voracious: seeking out new sources, investigating them, and evaluating them using their background beliefs. They investigate the reliability of purported experts and discredit experts when they have apparently good reason to do so, using their background beliefs. Our teenager seems, in fact, to be behaving with many epistemic virtues. They are not at all lazy; they are proactive in seeking out new sources. They are not blindly trusting; they investigate claims of epistemic authority and decide for themselves, using all the evidence and beliefs that they presently accept, whether to accept or deny the purported expertise of others. They have theories, which they have acquired by reasonable methods, predicting the maliciousness of outsiders; they increase their trust in those theories when their predictions are confirmed.¹⁰

The worry here is that agents raised within an echo chamber are, through no fault of their own, epistemically trapped. Their earnest attempts at good epistemic practices are transformed into something epistemically harmful by the social structure into which they have been embedded and which they have ingested. Paul Smart has argued for the possibility of a transformative social epistemic phenomenon which he dubs “mandevillian intelligence”, in honor of Bernard Mandeville. Mandeville argued that, in the right social context, individual vices could lead to collective economic prosperity. For a certain kind of economic theorist, capitalism is such a transformative structure: individuals act selfishly, but the structure of the market transforms that selfishness into virtuous collective action. According to Smart, there is an epistemic analog: the mandevillian intelligence, which transforms the individual epistemic vices of its members into a collective epistemic virtue by virtue of the social structure into which they are embedded (Smart 2017). Intellectual stubbornness, for example, might be an intellectual vice for individuals. But set those stubborn individuals in a properly arranged social structure (like, perhaps, academia) and you might get a collective system that properly explores every relevant nook and cranny with optimal thoroughness. But echo chambers are the very opposite; they are reverse-mandevillian intelligences. Echo chambers are social epistemic structures which convert individually epistemically virtuous activity into collective epistemic vice. In fact, the reverse-mandevillian nature contributes to the stickiness of the echo chamber trap. If our teenager self-reflects on their epistemic practices, what they will see might be rather gratifying. Their epistemic behavior might very well be earnest, vigorous, and engaged. It is their external context – the social epistemic system into which they have been unluckily raised – which makes such behavior problematic.

Contrast this account with Quassim Cassam’s treatment of Oliver, his fictional 9/11 conspiracy theorist. Oliver believes that the collapse of the twin towers was an inside job, and he is happy to provide reasons and point to supporting evidence from a great many conspiracy theorist websites. Says Cassam: Oliver is obviously mistaken – Oliver relies on outrageous, baseless claims from clearly discredited sources. The best explanation for Oliver’s beliefs is in terms of epistemic vice – that is, in terms of Oliver’s bad intellectual character traits. Oliver is “gullible, cynical, and prejudiced”, says Cassam. Oliver is gullible with regard to his conspiracy theorist sites, cynical with regard to the mainstream media, and his prejudice consists of, among other things, intellectual pride, wishful

¹⁰ For a parallel argument about the invidiousness of background prejudicial beliefs, see Begby (2013).

thinking, closed-mindedness, and a lack of thoroughness (Cassam 2016: 162–4). And I certainly grant that such epistemic vices can lead to these sorts of beliefs. But the story of our hapless teenager offers an alternate epistemic path to such beliefs and such narrow-casted trust – one in which epistemically virtuous character traits have been wrong-footed by the social epistemic structure in which the agent has been embedded. The crucial difference between the reverse-mandevillian account and Cassam’s account is where the brunt of the responsibility lies. In Cassam’s account, the responsibility lies with the individual, and their own intellectual habits and practices.¹¹ In a reverse-mandevillian account, a significant part of the responsibility lies with the social structure in which the actors are embedded. The epistemic vice is a feature of the collective intelligence, rather than of the individual. Or, if one is averse to thinking in terms of collective intelligences, here’s a conceptually minimal way to put the claim: echo chambers are local background conditions that turn generally good epistemic practices into locally unreliable ones.

But the possibility of a truly faultless epistemic agent, wholly misled by an echo chamber, also depends on the lack of an accessible escape route. So: are there escape routes from an echo chamber, and how reasonable is it to expect echo chamber members to discover and make use of them?

Here is one possible escape route. Consider Thomas Kelly’s discussion of belief polarization. Belief polarization is the tendency of individuals, once they believe that p , to increase their belief that p . Kelly argues that belief polarization works by the mechanism that, once an agent has acquired a belief, they tend to subject counter-arguments to greater scrutiny and give supporting arguments a relative pass. Thus, their critical reflection is likely to reinforce previously held beliefs. Kelly notes that the belief polarization violates the Commutativity of Evidence Principle:

The Commutativity of Evidence Principle: to the extent that what it is reasonable for one to believe depends on one’s total evidence, historical facts about the order in which that evidence is acquired make no difference to what it is reasonable for one to believe. (Kelly 2008: 616)

In short, belief polarization makes it matter what order they received the evidence, but the historical ordering of evidence ought not matter. Note that our epistemically hapless teenager has also violated the Commutativity of Evidence Principle. For them, the order in which they received the evidence matters very much. If they had been raised outside the echo chamber and fed a broader diet of epistemic sources before encountering the echo chamber, then they would likely have found the echo chamber’s world-view to be problematic. But since our teenager encountered the echo chamber and assimilated its beliefs first, their use of background beliefs to vet new sources will lead them to continually increase their trust in the echo chamber and their distrust of outsiders. Even if our echo chambered teenager eventually came to encounter all the same evidence as their epistemically free-range counterpart, their early education within the echo chamber would still be

¹¹ Note, however, that Cassam distinguishes between epistemic responsibility and epistemic blameworthiness, and does not take blameworthiness to necessarily follow from responsibility (Cassam 2016: 168–9). Cassam leaves room for the view that the individual’s intellectual vices were epistemically responsible for their bad beliefs, but that the individual wasn’t blameworthy for those vices, because the vices were inculcated in them at an early age. However, my complaint still stands, for I contest Cassam’s claim that the responsibility is in the individual.

decisive. So long as each new piece of evidence is assessed using the currently held set of beliefs, then early education in an echo chamber becomes domineeringly powerful.

However, the Commutativity of Evidence Principle suggests a way out. In order to free themselves of the echo chamber's grip, our teenager needs to undo the influence of the historical ordering of their encounters with the evidence. How could they possibly do this? Our teenager would have to suspend belief in all their particular background knowledge and restart the knowledge-gathering process, treating all testimony as equally viable. They would need to, in a sense, throw away all their beliefs and start over again. This suggests a process that, in its outlines, might sound awfully familiar. Our escape route turns out to be something like a modified version of Descartes' infamous method.

What proceeds from this point is admittedly something of a fantasy, but perhaps it is a fantasy from which we can eventually draw some sort of moral. The story of the history of Western epistemology might be cartoonishly summarized thusly: Descartes had a dream of radical intellectual autonomy. By his accounting, he came to realize that many of the beliefs he had acquired in his early life were false, and that those early false beliefs might have infected any number of other beliefs. His response was that famed method: to get rid of his beliefs and start over again, trusting no-one and nothing and only permitting those beliefs of which he was entirely certain. Call this the *Cartesian epistemic reboot*. But if recent epistemology has taught us anything, it's that this total reboot is nothing but a pipe dream. Any sort of reasonable epistemic life is essentially impossible without trusting the testimony of others (Hardwig 1985, 1991; Burge 1993; Faulkner 2000; Goldberg 2010; Zagzebski 2012).

But recall that the reason Descartes wanted to discard everything and start over from scratch – the motivation for his project, and not the method – was explained in the very first line of “Meditation 1”: He was worried by the falsehoods he had learned in childhood and the shakiness of the edifice that had been built from those falsehoods (Descartes 1984: 24). Our teenager faces a problem quite similar in structure. The credentialing structure of their upbringing is flawed; that credentialing structure has influenced any number of their other beliefs, and the degree of that influence is impossible to track. Furthermore, these later beliefs, approved by the echo chambers' credentialed sources, will often serve to reinforce that credentialing structure. The pernicious effect of an echo chamber cannot be attacked one belief at a time. Any single belief that our teenager re-considered would come under the influence of the network of the flawed background beliefs that sustains an echo chamber. What they need is some way to start over. In order to undo the influence of historical ordering, an epistemic agent will have to temporarily suspend belief in all their beliefs, *in particular their credentialing beliefs*, and start from scratch. But when they start from scratch, they need not disregard the testimony of others, nor need they hold to Descartes' stringent demand for certainty. Let's call this procedure the *social epistemic reboot*. In the social epistemic reboot, the agent is permitted, during the belief re-acquisition process, to trust that things are as they seem and to trust in the testimony of others. But they must begin afresh *socially*, by re-considering all testimonial sources with presumptive equanimity, without deploying their previous credentialing beliefs. Furthermore, they must discard all their other background beliefs, because those potentially arose from the flawed credential structure of the echo chamber, and very likely have been designed to support and reinforce that very credential structure. Our rebooter must take on the social epistemic posture that we might expect of a cognitive newborn: one of tentative, but defeasible, trust in all

apparent testimonial sources (Burge 1993; Nguyen 2011). This method will, if successfully applied, undo the historical dependence of our epistemic agent and remove the undue influence of the echo chamber. The social epistemic reboot is, theoretically at least, the escape route we've been searching for.¹²

This reboot, described in such clinical terms, might seem rather fantastical. But it is not, I think, utterly unrealistic. Consider the stories of actual escapees from echo chambers. Take, for example, the story of Derek Black, who was raised by a neo-Nazi father, groomed from childhood to be a neo-Nazi leader, and who became a teenaged breakout star of white nationalist talk radio. When Black left the movement, he went through a years-long process of self-transformation. He had to completely abandon his belief system, and he spent years re-building a world-view of his own, immersing himself broadly and open-mindedly in everything he'd missed – pop culture, Arabic literature, the pronouncements of the mainstream media and the US government, rap – all with an overall attitude of trust (Saslow 2016).

Of course, all we have shown so far is that the social epistemic reboot would, if pulled off, undo the effects of an echo chambered upbringing. Whether or not an epistemic agent might reasonably be expected to reboot, or blameworthy for failing to reboot, is a separate and significantly more difficult question. First, a social epistemic reboot might be psychologically impossible, or at least beyond what we could reasonably expect of normal epistemic agents. Second, what reason would an epistemic agent have to undertake a social epistemic reboot? Such an undertaking would be justified only if the agent had a significant reason to think that their belief system was systematically flawed. But echo chamber members don't seem likely to have access to any such apparent reason. After all, they have clear and coherent explanations for all the evidence and testimony they encounter. If this is all right then we arrive at a worrying conclusion: that echo chambers may, theoretically, be escapable, but we have little reason to expect members of echo chambers to actually realize that they are members of something that needs escaping.

What hope could we have, then, of motivating a reboot? Derek Black's own story gives us a hint. Black went to college and was shunned by almost everyone in his college community. But then Matthew Stevenson, a Jewish fellow undergraduate, began to invite Black to his Shabbat dinners. Stevenson was unfailingly kind, open, and generous, and he slowly earned Black's trust. This eventually led to a massive upheaval for Black – a slow dawning realization of the depths to which he had been systematically misled. Black went through a profound transformation and is now an anti-Nazi spokesperson.

The turning point seems to be precisely that Stevenson, an outsider, gained Black's trust. And this is exactly where we should expect the turning point to be. Since echo chambers work by building distrust towards outside members, then the route to unmaking them should involve cultivating trust between echo chamber members and outsiders. In order to motivate the social epistemic reboot, an echo chamber member needs to become aware of how much they are in the echo chamber's grip, and forming a trust relationship with an outsider might mediate that awareness. But how that trust could be reliably cultivated is a very difficult matter, and a topic for future investigation. We have, however, arrived at a tentative moral of the story. Echo chambers work by a manipulation of

¹² Note that the social epistemic reboot would also undo the effects of Begby's evidential pre-emption, since that pre-emption also depends on the historical ordering of received data.

trust. Thus, the route to undoing their influence is not through direct exposure to supposedly neutral facts and information; those sources have been pre-emptively undermined. It is to address the structures of discredit – to work to repair the broken trust between echo chamber members and the outside social world.¹³

REFERENCES

- Alcoff, L. M. 2007. 'Epistemologies of Ignorance: Three Types.' In S. Sullivan and N. Tuana (eds), *Race and Epistemologies of Ignorance*. Albany, NY: SUNY Press.
- An, J., Quercia, D. and Crowcroft, J. 2014. 'Partisan Sharing: Facebook Evidence and Societal Consequences.' In *Proceedings of the Second ACM Conference on Online Social Networks*. <http://cosn.acm.org/2014/files/cosno33f-anAembTS.pdf>.
- Baumgaertner, B. 2014. 'Yes, No, Maybe So: A Veritistic Approach to Echo Chambers Using a Trichotomous Belief Model.' *Synthese*, 191(11): 2549–69.
- Begby, E. 2013. 'The Epistemology of Prejudice.' *Thought: A Journal of Philosophy*, 2(2): 90–9.
- 2017. 'Evidential Pre-emption.' Presented at the American Philosophical Association Pacific Division 2017 Meeting.
- Bishop, B. 2009. *The Big Sort: Why the Clustering of Like-minded America is Tearing Us Apart*. Boston, MA: Houghton Mifflin Harcourt.
- Burge, T. 1993. 'Content Preservation.' *Philosophical Review*, 102(4): 457–88.
- Cassam, Q. 2016. 'Vice Epistemology.' *The Monist*, 99(2): 159.
- Coady, D. 2006. 'When Experts Disagree.' *Episteme*, 3(1–2): 68–79.
- 2012. *What to Believe Now*. Chichester: Wiley-Blackwell.
- Coates, D. D. 2010. 'Post-involvement Difficulties Experienced by Former Members of Charismatic Groups.' *Journal of Religious Health*, 49: 296–310.
- Corner, A., Whitmarsh, L. and Xenias, D. 2012. 'Uncertainty, Scepticism and Attitudes Towards Climate Change: Biased Assimilation and Attitude Polarisation.' *Climatic Change*, 114(3–4): 463–78.
- Dentith, M. R. X. 2015. 'When Inferring to a Conspiracy Theory Might be the Best Explanation.' *Social Epistemology*, 16: 572–91.
- 2017. 'The Problem of Conspiracy.' *Argumenta*, 5: 1–16.
- Descartes, R. 1984. *The Philosophical Writings of Descartes, Vol. 2*. Translated by J. Cottingham, R. Stoothoff and D. Murdoch. Cambridge: Cambridge University Press.
- El-Bermawy, M. 2016. 'Your Filter Bubble is Destroying Democracy.' <https://www.wired.com/2016/11/filter-bubble-destroying-democracy/>.
- Faulkner, P. 2000. 'The Social Character of Testimonial Knowledge.' *Journal of Philosophy*, 97(11): 581–601.
- Flaxman, S., Goel, S. and Rao, J. M. 2016. 'Filter Bubbles, Echo Chambers, and Online News Consumption.' *Public Opinion Quarterly*, 80: 298–320.
- Fricker, M. 2011. 'Rational Authority and Social Power: Towards a Truly Social Epistemology.' In A. Goldman (ed.), *Social Epistemology – Essential Readings*, pp. 54–70. New York, NY: Oxford University Press.
- Gilbert, E., Bergstrom, T. and Karahalios, K. 2009. 'Blogs are Echo Chambers: Blogs are Echo Chambers.' In *42nd Hawaii International Conference on System Sciences*. doi: 10.1109/HICSS.2009.91
- Goldberg, S. 2010. *Relying on Others: An Essay in Epistemology*. Oxford: Oxford University Press.
- 2011. 'If That Were True I Would Have Heard About It By Now.' In *The Oxford Handbook of Social Epistemology*, pp. 92–108. Oxford: Oxford University Press.

13 I'd like to thank Kara Barnette, Endre Begby, Anthony Cross, Melissa Hughs, Eric Stencil, Matt Strohl, Shannon Mussett, Bekka Williams, the anonymous reviewers, and many others for their assistance with this paper.

- 2013. 'Epistemic Dependence in Testimonial Belief, in the Classroom and Beyond.' *Journal of Philosophy of Education*, 47(2): 168–86.
- Goldman, A. I. 2001. 'Experts: Which Ones Should You Trust?' *Philosophy and Phenomenological Research*, 63: 85–110.
- Hardwig, J. 1985. 'Epistemic Dependence.' *Journal of Philosophy*, 82(7): 335–49.
- 1991. 'The Role of Trust in Knowledge.' *Journal of Philosophy*, 88(12): 693–708.
- Jamieson, K. H. and Cappella, J. N. 2008. *Echo Chamber: Rush Limbaugh and the Conservative Media Establishment*. Oxford: Oxford University Press.
- Kahan, D. M. and Braman, D. 2006. 'Cultural Cognition and Public Policy.' *Yale Law & Policy Review*, 24: 149–72.
- Kelly, T. 2008. 'Disagreement, Dogmatism, and Belief Polarization.' *Journal of Philosophy*, 105(10): 611–33.
- Kitcher, P. 1993. *The Advancement of Science*. New York, NY: Oxford University Press.
- Lackey, J. 2013. 'Disagreement and Belief Dependence: Why Numbers Matter.' In J. Lackey and D. Christensen (eds), *The Epistemology of Disagreement: New Essays*, pp. 243–68. Oxford: Oxford University Press.
- Langone, M. D. 1994. 'Reflections on Post-Cult Recovery.' <http://www.icsahome.com/articles/reflections-on-post-cult-recovery-langone>.
- Lifton, R. 1991. 'Cult Formation.' *Cultic Studies Journal*, 8: 1–6.
- Mengus, B. 2016. 'Pizzagaters Aren't Giving This Shit Up.' *Gizmodo*. <http://gizmodo.com/pizzagaters-arent-giving-this-shit-up-1789692422>.
- Miller, B. and Record, I. 2013. 'Justified Belief in a Digital Age: On the Epistemic Implications of Secret Internet Technologies.' *Episteme*, 10(2): 117–34.
- Millgram, E. 2015. *The Great Endarkenment: Philosophy for An Age of Hyperspecialization*. Oxford: Oxford University Press.
- Mills, C. 2007. 'White Ignorance.' In *Race and Epistemologies of Ignorance*, pp. 11–38. Albany, NY: State University of New York Press.
- Nelson, J. L. and Webster, J. G. 2017. 'The Myth of Partisan Selective Exposure: A Portrait of the Online Political News Audience.' *Social Media + Society*, July–September: 1–13.
- Nguyen, C. T. 2010. 'Autonomy, Understanding, and Moral Disagreement.' *Philosophical Topics*, 38(2): 111–129.
- Nguyen, C. Thi. 2011. 'An Ethics of Uncertainty: Moral Disagreement and Moral Humility.' PhD dissertation. University of California at Los Angeles. Retrieved from ProQuest Dissertations and Theses database (UMI No. 3532448).
- 2018a. 'Cognitive Islands and Runaway Echo Chambers: Problems for Expert Dependence.' *Synthese*. <https://doi.org/10.1007/s11229-018-1692-0>.
- 2018b. 'Escape the Echo Chamber.' *Aeon Magazine*. <https://aeon.co/essays/why-its-as-hard-to-escape-an-echo-chamber-as-it-is-to-flee-a-cult>.
- 2018c. 'Expertise and the Fragmentation of Intellectual Autonomy.' *Philosophical Inquiries*, 6(2): 107–24.
- Orwell, G. 1968. 'Politics and the English Language.' In S. Orwell and I. Angos (eds), *The Collected Essays, Journalism, and Letters of George Orwell, Vol. 4*, 1st edition, pp. 121–46. New York, NY: Harcourt, Brace, Jovanovich.
- Pariser, E. 2011. *The Filter Bubble: What the Internet Is Hiding From You*. London: Penguin UK.
- Pollak, J. 2017. '#FakeNews: Mainstream Media Continue to Slander Breitbart.' *Breitbart*, 30 January. <http://www.breitbart.com/big-journalism/2017/01/30/fakenews-mainstream-media-continue-slander-breitbart/>.
- Robbins, T. and Anthony, D. 1982. 'Deprogramming, Brainwashing and the Medicalization of Deviant Religious Groups.' *Social Problems*, 29(3): 283–97.
- Robson, J. 2014. 'A Social Epistemology of Aesthetics: Belief Polarization, Echo Chambers and Aesthetic Judgement.' *Synthese*, 191(11): 2513–28.
- Saez-Trumper, D., Castillo, C. and Lalmas, M. 2013. 'Social Media News Communities: Gatekeeping, Coverage, and Statement Bias.' In *Proceedings of the 22nd ACM International Conference on Information & Knowledge Management*. <https://dl.acm.org/citation.cfm?id=2505623>.

- Saslow, E. 2016. 'The White Flight of Derek Black.' *The Washington Post*, 15 October. https://www.washingtonpost.com/national/the-white-flight-of-derek-black/2016/10/15/ed5f906a-8f3b-11e6-a6a3-d50061aa9fae_story.html?utm_term=.67fooca25ce9.
- Singer, M. T. 1979. 'Coming Out of the Cults.' *Psychology Today*, 12(8): 72–82.
- Smart, P. R. 2017. 'Mandevillian Intelligence.' *Synthese*. <https://doi.org/10.1007/s11229-017-1414-z>.
- Sunstein, C. R. 2001. *Republic.com*. Princeton, NJ: Princeton University Press.
- 2009a. *Going to Extremes: How Like Minds Unite and Divide*. Oxford: Oxford University Press.
- 2009b. *Republic.com 2.0*. Princeton, NJ: Princeton University Press.
- Trouche, E., Johansson, P., Hall, L. and Mercier, H. 2016. 'The Selective Laziness of Reasoning.' *Cognitive Science*, 40(8): 2122–36.
- Vogt, P. J. and Goldman, A. 2016. 'Voyage Into Pizzagate.' <https://gimletmedia.com/episode/83-voyage-into-pizzagate/>.
- Watson, J. C. 2015. 'Filter Bubbles and the Public Use of Reason: Applying Epistemology to the Newsfeed.' In F. Scalabrino (ed.), *Social Epistemology and Technology: Toward Public Self-Awareness Regarding Technological Mediation*, pp. 47–58. London: Rowman & Littlefield.
- Weathers, C. 2014. 'CrossFit Is a Cult: Why So Many of Its Defenders Are So Defensive.' *Salon*. https://www.salon.com/2014/10/22/crossfit_is_a_cult_why_so_many_of_its_defenders_are_so_defensive_partner/.
- Wittgenstein, L. 2010. *Philosophical Investigations*. Hoboken, NJ: Wiley.
- Zagzebski, L. 2012. *Epistemic Authority*. New York, NY: Oxford University Press.

C. THI NGUYEN is Associate Professor of Philosophy at Utah Valley University. His research includes social epistemology, practical reasoning, and aesthetics; he is particularly interested in the ways in which reasoning and agency are socially embedded. His book, *Games: Agency as Art*, is forthcoming from Oxford University Press.
