

ARTICLE

Echo Chambers and Audio Signal Processing

Benjamin Elzinga* 

Georgetown University, Washington DC, USA

*Corresponding author. Email: be72@georgetown.edu

(Received 7 November 2019; revised 31 May 2020; accepted 15 July 2020;
first published online 30 September 2020)

Abstract

Following Cass Sunstein's popular treatment of the concept, echo chambers are often defined as environments which exclude contrary opinions through omission. C. Thi Nguyen contests the popular usage and defines echo chambers in terms of in-group trust and out-group distrust. In this paper, I argue for a more comprehensive treatment. While both exclusion by omission and out-group distrust help sustain echo chambers, neither defines the phenomenon. I develop a social network model of echo chambers which focuses on the role of belief-reinforcing echoes. First, I argue that the model allows us to incorporate Nguyen's main point about distrust without construing other commentators as deeply mistaken about the nature of echo chambers. Second, I use the model to develop an account of *collaborative resistance* and use it to clarify the role echo chambers play in spreading misinformation.

Keywords: Echo chamber; misinformation; filter; fake news; feedback loop; social network; social media

1. Introduction

An *acoustic echo chamber* is a hollow enclosure used to create a reverberation effect for musical purposes. At one end of the echo chamber, a speaker produces a recording of someone singing, playing an instrument, or some other audio signal in need of additional processing. The sound from the speaker bounces off the walls repeatedly, creating a cascade of echoes until it eventually decays. A microphone picks up the auditory cascade and feeds it into another audio recorder, thereby preserving a modified version of the original recording. The term typically refers to rooms designed with this specific purpose in mind, but the reverberation effect also occurs in cathedrals, caves, and other hollow spaces which echo sound back to you. A *metaphorical echo chamber* is a social environment which echoes your own viewpoint back to you.

The events surrounding the 2016 presidential election in the U.S. sparked renewed interest in the problem of echo chambers (see, e.g., Chater 2016; Emba 2016; Hooton 2016; Hosanagar 2016; Khalid 2017). Taking inspiration from Cass Sunstein's earlier treatment of concept, the major focus was on how social media platforms help users keep opposing viewpoints out of their newsfeeds (Sunstein 2001, 2009a, 2009b; see also Pariser 2011). As a result, "echo chamber" in current discussions usually refers to an environment which excludes contrary opinions through omission. In a recent treatment of the concept, C. Thi Nguyen contests popular usage and defines echo chambers in terms of out-group distrust (Nguyen 2020). In this paper, I argue for a

more comprehensive treatment. An echo chamber is a social network where beliefs are robustly and routinely reinforced through the echoing of consonant belief expressions throughout the network. While both exclusion and distrust may play a role in creating and sustaining echo chambers, neither defines the phenomenon.

My treatment focuses on two deeply related questions. The first question is “what’s an echo chamber?” My goal is to reconcile Sunstein and Nguyen’s answers to the question while otherwise sticking as closely to the echo chamber metaphor as possible. Both Sunstein and Nguyen capture something important about the phenomena, but their accounts are in tension with one another. My aim is to resolve the tension without completely undermining either contribution. I begin by introducing a basic social network model of echo chambers (section 2) and enrich the model further by exploring the role of epistemic filters (section 3). While both filters of omission and distrust play an ineliminable causal role in creating and sustaining echo chambers within pluralistic societies, I argue that such filters are not constitutive features of echo chambers themselves (section 4).

The second question is “is an echo chamber as such a problematic epistemic environment?” Most commentators focus on the role echo chambers play in spreading false beliefs or producing other negative epistemic effects. In *Network Propaganda*, Yochai Benkler, Robert Faris, and Hal Roberts offer a comprehensive treatment of the origins of misinformation focusing on the 2016 U.S. election (Benkler *et al.* 2018). I use their account not only to illustrate my social network model of echo chambers but also to raise important questions about their explanatory role (section 5). I argue that the most salient cause of false beliefs within a community is the absence of mechanisms which constrain the echoing of consonant belief expressions to the world of fact (section 6). Echo chambers, on the other hand, are best suited for explaining the resilience of communities built around false beliefs and their resistance to outside interventions and internal dissent (section 7). Finally, with a better understanding of what is worrisome about problematic echo chambers in hand, I explore potential solutions (section 8).

2. Social networks

In his influential treatment, Cass Sunstein emphasizes *personal echo chambers* and our tendency to selectively curate our social networks. My treatment instead focuses on the social networks themselves. This is after all what creates the belief-reinforcing echoes and therefore constitutes the echo chamber. Since an echo chamber is a social network, the idea is usefully represented using a *social network diagram*. Social network diagrams consist in *nodes* representing individuals or groups and *edges* or *links* representing the connections between nodes. For example, nodes might represent individual Facebook users, and links might represent friendships as defined by the platform. Social network diagrams are typically generated out of raw data in order to represent social networks in an accessible format. While I intend to borrow design elements from social network diagrams, my purpose is fundamentally different. My goal is not to represent any particular existing social network but to provide a general model for a diverse range of possible social relationships and processes. Later I will explore the ways in which these processes might occur in the real world (section 5) and argue that we all in fact occupy echo chambers (section 6). Since my primary goal is to illustrate a concept, I don’t take a stand here on whether the modified social network diagram I construct would be a useful tool for data analysis (see, e.g., Benkler *et al.* 2018) or determining optimal arrangements of information flow (see, e.g., O’Connor and Weatherall 2019).¹

¹Since the final model I settle on represents a great number of agent properties, many of which would be hard to measure, I suspect it wouldn’t be particularly useful for the former.




Figure 1. A partial representation of a social network in action.

In the echo chamber diagram (Figure 1), nodes represent epistemic agents in a broad sense of the term. An agent is an individual, group, or institution capable of processing and producing *belief expressions*. Belief expressions include the transmission of any propositionally structured representation and especially cases of *telling*, where one agent gives their assurance to another agent that something is true (Moran 2005).² I assume for the purposes of the model that individual agents routinely seek out belief-confirming information (compare Benkler *et al.* 2018: 76).³ Gray arrows represent the possible channels for belief transmission while black arrows represent typical patterns of information flow. Suppose that news outlet E publishes a story or article which is distributed to B and D and reinforces some of their beliefs. Those agents then transmit the news story or related contents by posting a link to the news article on social media or through conversation. Some nodes, like A and G, would not be made aware of the story or its contents otherwise, while others like B and D receive the news through multiple channels. To appreciate the full force of the idea, it's helpful to imagine the other possible connections between the nodes represented in the diagram and other agents omitted from the model. The components of the diagram can be recombined to represent an infinite number of possible social arrangements involving a diverse range of agent-types and complex connections.

While reliable channels for information flow are necessary for ideas to spread, an environment supporting belief transmission and reinforcement is not necessarily an echo chamber. An echo chamber is an environment which transmits and reinforces beliefs members already hold or at least beliefs consonant with their shared perspective. For example, a politically liberal chamber would be one where members have significant overlap in left-leaning beliefs. A conservative echo chamber would be one where members have significant overlap in right-leaning beliefs. What distinguishes different echo chambers is therefore the particular contents of the echoing beliefs, but what makes a social network count as an echo chamber is the fact that the shared beliefs of members are echoed back to them. In the current model, grey edges represent the available

²This is a stipulative definition of “belief expression” tailored towards the aims of this essay. Notice that on this usage someone can produce a belief expression that P even if they don't believe P themselves.

³That's not to say that agents never seek out or accept falsifying evidence or that the tendency to seek confirmation is always epistemically vicious (i.e. a form of wishful thinking or confirmation bias). Agents may seek belief confirmation simply because they want their belief to be true, but they may also seek it out because they think, perhaps correctly, that the belief is reasonable and desire further confirmation.

information channels within a social environment for the transmission of any belief, but in any epistemically pluralistic society there will be a diverse range of views making use of those channels. This suggests that we also need to add the idea that agents selectively make use of certain channels and give differential uptake to information flows depending on their content and source. I will nevertheless qualify this claim in section 4 and argue that selective uptake is not necessary for echo chambers strictly speaking.

3. Filters

Agents give differential uptake to belief expressions by imposing various filter mechanisms.⁴ To keep with the precedent set by the echo chamber metaphor, consider a musical example. A “low-pass” filter is a mechanism which lets audio frequencies below a user-specified cutoff pass through while attenuating the amplitude of higher frequencies. A “high-pass” filter does the exact opposite. There are a variety of other musical filters which all do the same sort of thing: let in some frequencies and exclude others. In his treatment of echo chambers, Nguyen identifies two forms of social epistemic filtering. The first is exclusion by omission, which creates what Nguyen calls “epistemic bubbles”.⁵ Omission occurs via “selective exposure” and what I’ll call “third-party filtering” (Nguyen 2020: 143–4). Selective exposure refers to an individual’s tendency to seek out, engage, or consume information from like-minded sources. People tend to befriend individuals with similar political, scientific, or religious viewpoints or with similar world-views generally. Importantly in this context, this means that people tend to associate with those who share many of the same beliefs, knowledge-seeking skills, and attitudes.

Third-party filtering refers to the ways in which other agents or organizations shape, constrain, or modify what information is available to an agent. A major focus in the recent literature on misinformation and social media is on the role of personalization algorithms. By predicting what one would like to see in one’s news feed or search results based on one’s past interactions, websites like Facebook and Google supposedly create what Eli Pariser calls a “filter bubble” (Pariser 2011). For example, if you tend to search for liberal news sources, these algorithms supposedly ensure that future searches will be more likely to yield results consonant with your liberal beliefs. Although there is some debate over the extent of the effect, the online filter bubble would be a case of other agents excluding certain information from one’s environment rather than self-directed selective exposure (see Bakshy *et al.* 2015; Pariser 2015; Benkler *et al.* 2018: 289–91).

Following Sunstein (Sunstein 2001, 2009a, 2009b). Nguyen, on the other hand, uses the term “echo chamber” to refer exclusively to an “epistemic community which creates a significant disparity in trust between members and non-members” (Nguyen 2020: 146). Such echo chambers are created through “epistemic discrediting” of outsiders and “epistemic credential amplifying” of insiders where in-group membership requires adherence to a core set of beliefs involving beliefs that support the disparity in trust (Nguyen 2020: 146). Exclusion through credibility distribution is distinct from mere omission. To

⁴Throughout the article, I will use “mechanism” to refer to any organized system of parts which reliably produces some specific result (compare Craver and Tabery 2017).

⁵Nguyen defines an epistemic bubble as a “social epistemic structure which has inadequate coverage through a process of exclusion by omission” (Nguyen 2020: 143). Notice that Nguyen adds the idea that bubbles have “inadequate coverage”. This suggests that bubbles necessarily tend to produce bad epistemic effects (e.g., false beliefs). But notice that nothing about exclusion by omission in itself implies bad epistemic effects. It’s perfectly possible to filter out false beliefs and let true ones through. Similar issues will arise in section 6, but my focus there will be on distrust and echo chambers.

discredit someone or some group is to treat them as untrustworthy due to “unreliability, epistemic maliciousness, or dishonesty” or some other kind of “epistemic demerit” (Nguyen 2020: 146). While outgroup members are assigned low credibility, ingroup members are assigned high levels of trust. Nguyen further points out that in epistemic bubbles individuals simply aren’t exposed to alternative viewpoints, but in echo chambers they typically are (see also Garrett *et al.* 2016). They may in fact even voraciously or obsessively consume opposing sources of information. Since they assign them low credibility, however, those sources are unlikely to prompt belief revision.

Before evaluating Nguyen’s distinction between epistemic bubbles and echo chambers, let’s see how these ideas work with the social network model. Since both omission and distrust involve an element of exclusion or selective uptake to information, they both fit the filter metaphor. Agents filter out certain belief expressions through active omission or distrust. To enrich the social network model, we can use dashed circles surrounding a node to represent omission-filters and dotted circles within a node to represent credibility-filters. Gray links are removed for the sake of simplicity, and black links still represent typical patterns of information flow.

In the enriched model (Figure 2), node E on the right might be Fox News or some other paradigmatically right-leaning news source. On the left node, C might be MSNBC or some traditionally left-leaning news source. The other nodes might be members of the public on the political right, like nodes B and G, or left, like nodes A and F, and node D in the middle might be interpreted as on the political center or center-right. Nodes on the right and left are connected internally as with D and B, but they may be cut the other side through the use of various filters as with D and A. It’s important to keep in mind that the model doesn’t represent every possible node or connection between nodes. Agents on the left might also watch Fox News (but through a lens of distrust), and there will often be an exchange of information between news sources themselves. I’ll return to both ideas in section 5.

The enriched model has more resources for making sense of how echo chambers reinforce beliefs in a pluralistic society. While consonant belief expressions help




Figure 2. An enriched representation of a social network. Nodes in the center and on the right are cut off from nodes on the left by means of various filters. Belief expressions may be blocked (represented by black dots) by a subtractive filter or prevented from significantly impacting one’s belief in the expressed content by an all-pass filter (it may nevertheless have an effect on other beliefs, e.g., one’s beliefs about what a distrusted agent believes). Belief expressions which get through (represented by black arrowheads) may help support, reinforce, or modify the agent’s belief in the expressed content. Grey arrows representing available information channels are removed for the sake of simplicity, black arrows represent typical patterns of information flow, and filters are understood as applying only to incoming messages. The absence of a link between nodes is not intended to be significant.

support, reinforce, or modify an echo chamber member's beliefs, divergent opinions are kept out through active filters of omission or distrust. Although my focus is on political echo chambers, it's worth emphasizing that echo chambers may form around any topic whether substantive or trivial in import or broad or narrow in scope (see also Nguyen 2020: 149–50).

4. Mixed metaphors

Nguyen introduces a sharp distinction between “epistemic bubbles” and “echo chambers” because commentators often ignore the difference between exclusion by omission and exclusion through credibility distribution (Nguyen 2020: 152–3). He argues that failing to appreciate the distinction has led some to underestimate problematic echo chambers or propose ineffective solutions. It's relatively easy to include more views in one's environment and many agents do expose themselves to alternate viewpoints, but once the walls of distrust are up they can be difficult to tear down. I agree that the substance of Nguyen's distinction is important. An echo chamber reinforces common beliefs through the echoing of consonant beliefs. In an epistemically diverse community, both omission-filters and credibility-filters play an important role in the process of belief-reinforcement. When omission-filters are in place, one will encounter many belief-supporting echoes and few dissenting opinions. When credibility-filters are in place, the dissenting opinions that do get through won't have a significant impact because they will be interpreted through a filter of distrust. Both filter-types effectively remove potential obstacles between the echoing of consonant beliefs and belief reinforcement.

My view therefore allows us to incorporate Nguyen's main point without construing other commentators as deeply mistaken. Consider the following from Jennifer Lackey:

There are two components to an echo chamber – first, there is an opinion that is repeated and reinforced, and second, this occurs in an enclosed system or “chamber,” such as a social network, allowing the opinion to “echo.” Dissenting voices are either absent or drowned out, and the original opinion is amplified through re-sharing. (Lackey 2018)

If Nguyen's account is correct, Lackey is deeply confused. In trying to characterize echo chambers, she has mistakenly described epistemic bubbles. But this can't be quite right. Lackey has clearly latched onto the central idea behind the metaphor: in an echo chamber your own view echoes back to you. While the most problematic echo chambers may rely on credibility-filters, some echo chambers may rely solely on exclusion by omission.

Lackey is also correct in thinking that echo chambers, in principle, don't require filters of any kind. A filter is a mechanism for *active exclusion* by omission or distrust.⁶ There is a difference, for example, between a frequency simply being absent and using a filter to selectively remove a frequency. While we often actively omit or distrust conflicting viewpoints, an echo chamber where dissenting opinions are simply absent is possible. Imagine a community where there are no contrary points of view, the dissenting minority is too small to make an impact, or perhaps lacks the material means to spread their message. Imagine further that the beliefs of like-minded individuals within the majority are still robustly and routinely reinforced through belief-echoes. No filters are present because no filters are needed, but intuitively such a community still counts

⁶Compare Nguyen (2020: 146). Notice, however, that for Nguyen epistemic bubbles can form through mere omission without active exclusion (Nguyen 2020: 143). Exclusion by omission, however, is by definition an active form of exclusion.

as an echo chamber. It's just not a very complex one, and probably not one found in any pluralistic society.⁷ Therefore, while filters may be *causally necessary* to create and sustain echo chambers in certain circumstances, they are not *constitutively or conceptually necessary* features of echo chambers. Both Sunstein and Nguyen make the mistake of latching onto some typical feature of echo chambers (exclusion by omission and exclusion by distrust respectively) and taking it to be constitutive of the phenomena.

It may be that Nguyen's intention was to clear up some of the confusion in the literature on echo chambers and misinformation by stipulating distinct uses for common terms. If Nguyen was only interested in distinguishing between two phenomena, then there is no substantive disagreement between us since my account is perfectly compatible with his distinction. It's not clear, however, whether Nguyen takes himself to be merely stipulating usages for common terms. In marking out the distinction between bubbles and echo chambers and claiming that Sunstein only manages to capture the former, Nguyen suggests that Sunstein failed to understand what echo chambers are. Nguyen may have something like the following line of reasoning in mind. We use the term "echo chamber" to refer to social networks with a shared set of core beliefs with a few common, identifying features: (1) those beliefs are reinforced by that community's social epistemic properties or activities, (2) those beliefs are highly resistant to outsider intervention or correction, and (3) those beliefs tend to be false due to some epistemic defect inherent to echo chambers. To provide an account of echo chambers would be to identify the most salient, common cause which explains these core features. Nguyen argues that exclusion by omission fails to explain (2), the resilience of echo chambers in the face of outsider intervention, and that this feature, along with the other two, is explained by exclusion through distrust. Therefore, we should provide an account of echo chambers in terms of distrust.

It would be helpful if Nguyen clarified whether these were his starting assumptions because they raise a number of pressing questions.⁸ The first identifying feature, that network member beliefs are reinforced by that community's epistemic properties or activities, is uncontroversial within the literature and works well with the metaphor. However, intuitions may vary when it comes to the second feature, the idea that member beliefs are highly resistant to outsider intervention or correction. Throughout his discussion, Nguyen seems to have the worst echo chambers in mind. He focuses on extremist right-wingers (2020: 145–6, 150–3), climate change deniers (2020: 150–1), conspiracy theorists (2020: 148, 155), and neo-Nazis (2020: 158) and draws comparisons with cults and cult-indoctrination techniques (2020: 147). But focusing on the most problematic and highly resistant epistemic communities may distort one's sense of what echo chambers are. As I pointed out a moment ago, it's perfectly natural to use the concept to describe cases where dissenting opinions are simply absent and members are not highly resistant to criticism. My more encompassing definition allows these to count as echo chambers as well without ruling more extreme cases out.⁹

⁷But see work on epistemic injustice and oppression (see, e.g., Fricker 2009; Dotson 2014). Hermeneutical marginalization, for example, can create conditions in which conflicting viewpoints simply aren't formulated or expressed, even within nominally pluralistic societies.

⁸It is possible that Nguyen intended his definition to allow for good echo chambers and so he doesn't accept (3). As I argue in section 6, (3) doesn't follow from Nguyen's official definition of echo chambers. At the same time, he frequently makes strong claims about the tendency of echo chambers to produce negative effects. This makes it unclear whether he intended to capture (3) in his definition (but failed to do so) or leave echo chambers neutral while focusing on bad cases.

⁹Nguyen of course would describe the milder cases in terms of "epistemic bubbles". If he is merely stipulating uses for "epistemic bubble" and "echo chamber", this is largely unobjectionable. But the way he marks out his distinctions is not neutral, and this choice in terminology may even add new confusions

The third identifying feature, that echo chambers are built around false beliefs, should not be taken for granted either. Some philosophers maintain that echo chambers are in themselves neither good nor bad (Fantl 2018; Lackey 2018). Most commentators focus on the negative effects of echo chambers, but this doesn't directly imply anything about the value of echo chambers. Compare the term "conspiracy theory". While the term is almost always used to pick out false or poorly supported theories, it can also be used to refer to any theory offering an explanation in terms of a conspiracy. Some such theories are true and well-supported, like the conspiracy theory associated with the Watergate scandal, while others are baseless and false. So long as one makes it clear whether the negative or neutral sense is being deployed, adoption of either usage may be acceptable. Perhaps the same goes for echo chambers.

To explore the question of echo chamber neutrality further, it will help to examine some concrete cases where echo chambers play a role in bringing about bad epistemic effects. I will focus on the role of right-wing echo chambers in the context of the 2016 U.S. election (section 5) and on the flat-earth conspiracy (section 6).

5. Case study

The recent discussion about echo chambers focuses on their role in spreading misinformation through online social media platforms like Reddit, Facebook, Twitter, and so on. In *Network Propaganda*, Benkler *et al.* offer a comprehensive treatment of the origins of misinformation focusing on the 2016 U.S. election, but they argue against the currently dominant "the internet did it" narrative (Benkler *et al.* 2018: 292):

Although a growing proportion of the U.S. population uses Facebook and other social media as primary sources of news, a large portion, particularly those over 50 but even many who are younger, still rely on broadcast television and cable news. Talk radio remains an important source of information and ideology for many Americans. Although the number of people who read newspapers and news magazines is shrinking, it remains significant, and political news is also brought to our attention by many other means, including face-to-face conversations in offline social networks, email, campaign advertising, social media, and family dinners. (Benkler *et al.* 2018: 45)

Benkler *et al.* argue that the primary source of misinformation is the adoption of divergent media practices on the right, the origins of which precede the existence of social media platforms. Exploring their view will help illustrate the echo chamber model further and bring the relationship between echo chambers and the spread of misinformation into clearer focus.

By analyzing the flow of information across news media platforms, Benkler *et al.* argue that the contemporary media ecosystem is highly asymmetric. While the right is isolated from the rest of the media ecosystem, the left is integrated with center and center-right media outlets which adhere more closely to professional standards of journalism. This exposes the left to sources of meaningful disconfirmation not available on the right. To capture the idea in terms of the social network model, we can introduce a variation where nodes represent collectives (Figure 3). While the right puts up

and inconsistencies into the literature since the term "echo chamber" intuitively applies in cases of mere omission. The only explicit reason Nguyen gives for defining echo chambers in terms of distrust is that it fits the phenomena Jamieson and Cappella describe in *Echo Chamber: Rush Limbaugh and the Conservative Media Establishment* (Jamieson and Cappella 2008; Nguyen 2020: 145). But it's not clear why one should think that the phenomena Jamieson and Cappella describe are exhaustive of the forms echo chambers can take.

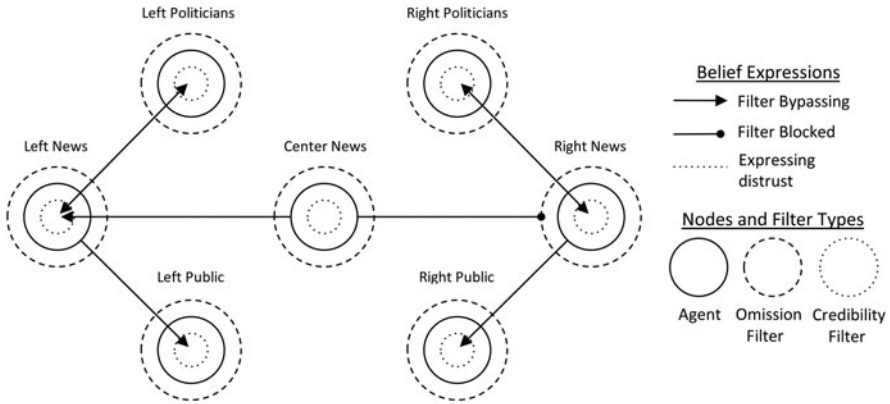


Figure 3. An asymmetric media network. While disconfirming news stories from the center regularly flow to the left, the right selectively blocks disconfirming stories coming from the center. In the case of news media institutions, a black arrowhead penetrating to the center of the node represents an expressed content being reported as fact (or minimally reported as a piece of valuable evidence or testimony supporting a fact). A black dot on a dashed circle represents a claim being blocked by subtractive filters which determine whether a story is news-worthy in accordance with prevailing journalistic norms. While any news institution will attempt to exclude irrelevant or uninteresting stories, part of Benkler *et al.*'s claim is that the right-wing media also actively filters out identity-disconfirming stories to a greater extent than left-wing media.

omission-filters, cutting themselves off from the center, information flows more freely from the center to the left. As a result, the right excludes potential sources of disconfirmation, and the left exposes itself to evidence or claims from credible sources which might undermine identity-confirming narratives. To illustrate the point, Benkler *et al.* (2018: 85–99) compare the relative success of sexual abuse stories during the 2016 election. While conspiratorial narratives linking the Clintons to child sex trafficking were picked up and widely disseminated by mainstream media on the right, similar stories alleging that Trump sexually abused minors failed to receive uptake by mainstream media of the left.

To account for the origins of insular right-wing media, Benkler *et al.* offer a subtle examination of U.S. news media history, but a brief summary puts emphasis on the rise of Rush Limbaugh, Fox News, and later Breitbart (Benkler *et al.* 2018: Ch. 10–11). Each adopts a strategy that emphasizes “partisan-confirming news over truth” and helps their audience reduce or nullify any discomfort they may feel in response to disconfirming evidence by “telling them that the outlets providing disconfirming news are not trustworthy” (Benkler *et al.* 2018: 78). Figure 4 shows how the latter may be represented. Finally, Benkler *et al.* argue that even when the original breeding grounds of misinformation are social media platforms, their successful distribution relies on a “propaganda pipeline” which takes ideas from the periphery to core mainstream news sources “through a series of well-known amplification sites, most prominently Infowars and [the Drudge Report]” (Benkler *et al.* 2018: 223) (Figure 5). While online propagandists regularly and strategically attempt to insert sensationalist memes into the broader public, not all attempts succeed. They rely on a willing mainstream media for success (Benkler *et al.* 2018: 225–33).

Although they acknowledge the role of individual patterns of omission and distrust as well as the amplificatory effects of new technology, Benkler *et al.* treat these as parts of a broader system where the most salient problem is the norms governing right-wing media practices. The result is a political ecosystem, involving members of the public,

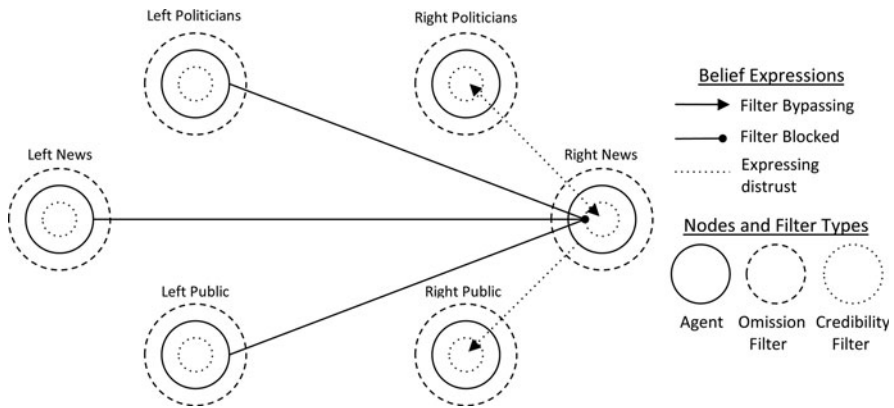


Figure 4. Right-wing distrust amplification: when ideas pass through right-wing subtractive filters, they are typically blocked by credibility filters. Even if the right-wing media reports on the left, they do so in a way that expresses the idea that the left is not to be trusted.

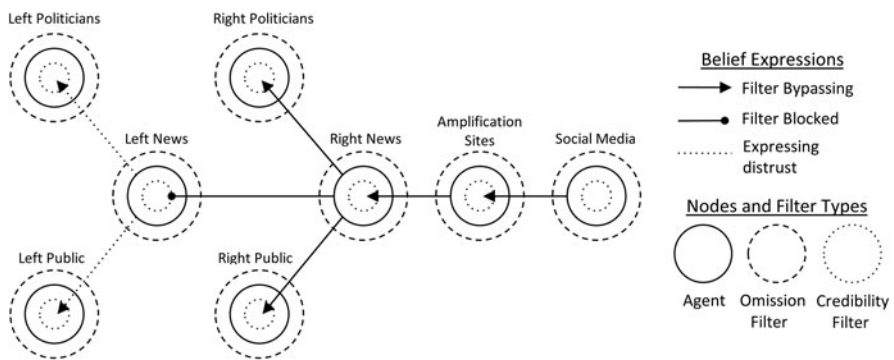


Figure 5. A representation of the ‘propaganda pipeline’. While information travels freely from peripheral online platforms to the right-wing echo chamber, leftwing media views the right through a lens of distrust.

news sources, political pundits, and politicians, who sanction departure from identity-confirming narratives more heavily than they sanction departure from the truth. Benkler *et al.*, therefore, explain the “epistemic crisis” gripping contemporary U.S. politics in terms of the echoing of misinformation supported by a combination of omission and credibility filters. Whether or not their diagnosis is correct is not of particular concern in the present context. Although they don’t clearly or unambiguously put the point this way, their interest is in echo chambers as I have defined them.

6. Same boat, different echo chamber?

An echo chamber is a social network where beliefs are robustly and routinely reinforced through the echoing of consonant belief expressions throughout the network. In an epistemically diverse society, echo chambers will invariably make use of various filters. But if this definition is correct, then we all (hermits aside) occupy echo chambers. In forming beliefs about matters outside of one’s immediate environment, one benefits not only from putting up various filters but also from the support of like-minded individuals

who gain support from each other as well. Consider the social network diagram representing the right/left divide. While nodes on the right give selective uptake to certain belief expressions, so do the nodes on the left. The right may be substantively cut off from the truth, but the definition of echo chambers doesn't describe one's relationship to the truth. It only describes one's relationship to other agents. This consequence may be cause for concern. In light of similar considerations, Jennifer Lackey warns that focusing on echo chambers may lead to a dangerous "we are all in the same boat" mentality (Lackey 2018).

To illustrate the worry in more extreme terms, consider the flat-earth movement. Flat-earthers claim that the land mass we inhabit is flat and that the prevailing round-earth theory is a lie perpetuated by the government and mainstream educational systems. Genuine flat-earthers, as opposed to online trolls and memers, occupy an echo chamber involving exclusion, distrust, and the repetition of confidence boosting ideas (Clark 2018). By the very same token, however, their critics and members of the scientifically literate public also occupy an echo chamber. Flat-earth views are excluded from mainstream education systems, most people distrust (and perhaps ridicule and mock) flat-earthers, and they gain confidence in the round-earth theory from the fact that just about everyone else agrees with them. If both conspiracy theorists and their debunkers occupy echo chambers, then it seems that they are in the same epistemic boat. Neither has a legitimate claim against the other that the other doesn't have against them, and the two echo chambers are on an equal epistemic footing.

While my view does imply that we all occupy echo chambers, it does not imply that we are all in the same boat. Echo chambers are in themselves neutral, but good and bad echo chambers are distinguished by the presence or absence of *truth-conducive mechanisms*. Such mechanisms open up reverberating beliefs to constraints from the world of fact. On a broadly reliabilist interpretation, a truth-conducive mechanism is an organized system of parts which tends to yield significantly more true beliefs than false beliefs. Since mechanisms produce processes, one may simply speak in the more traditional terms of "reliable belief-forming processes". In this context, reliability may be understood in a "frequency sense" referring to the success or failure of a process within the actual world (Goldman 1979: 11). Belief-forming processes of the relevant kind may take place in an individual's head, but they may also involve multiple agents, institutions, environmental features and so on. Truth-conducive mechanisms will plausibly include mechanisms for checking ideas coming into an echo chamber and mechanisms for double-checking ideas already inside where "checks" are potential sources of disconfirmation. In a good echo chamber, true beliefs will tend to echo more than false beliefs because truth-conducive mechanisms are in place, providing useful "epistemic friction" and exposing false beliefs to scrutiny.¹⁰

The problem with the flat-earth echo chamber is that the relevant belief-echoes are largely unconstrained by truth-conducive mechanisms. Their belief system is therefore subject to what John McDowell colorfully refers to as a "frictionless spinning in a void"

¹⁰I borrow the phrase "epistemic friction" from Medina (2013) who derives it from Wittgenstein. One may want to inquire into the nature of true-conducive mechanisms and how they provide genuine checks on our beliefs. There are two forms this inquiry might take. One may ask after the general features that any epistemic practice must include for the world to provide genuine checks on our beliefs. This is a question that epistemologists and philosophers of science have long been in the business of answering (for two very different angles, see Haugeland 2002; Medina 2013). One may also ask how to install and maintain those mechanisms in particular truth-seeking contexts. This is a question that scientists and policymakers (in a broad sense of the term) have long been in the business of answering (see, for example, section 8 and especially fn. 16). My reasons for leaving out a detailed answer to these questions is not because there are no viable answers but because there are too many answers to adequately address in a short article.

(McDowell 1996: 11). This suggests that the spread of misinformation is best explained by the presence of belief-forming mechanisms which are actively falsehood-conducive. Such mechanisms may include epistemically vicious cognitive biases and belief-forming practices (Battaly 2014; Cassam 2016), hyper-moralized and emotionally charged communication styles (Brady *et al.* 2017; Crockett 2017), ignorance and meta-ignorance (Dunning 2011; Medina 2013), unstable norms of online communication (Rini 2017), the way collective memories are shaped through conversation (Spinney 2017), social media and search engine algorithms (Pariser 2011), highly personalized and difficult to monitor political advertisements on the internet (Tufekci 2014) or, as the case study emphasizes, defective institutional practices (Benkler *et al.* 2018). Falsehood-conducive mechanisms may also be understood in a frequency sense. This means that the presence of such mechanisms contributes to the production of false beliefs significantly more often than true beliefs, supporting belief formation without or in direct abeyance of constraint by truth-conducive mechanisms.¹¹

If the cause of misinformation is in fact falsehood-conducive mechanisms, then one might think, as Lackey does, that echo chambers as such are not a problem at all (Lackey 2018). Echo chambers support both false belief systems and true belief systems, and so they aren't the cause of misinformation.¹² Instead, echo chambers are more or less global background conditions enabling the spread of consonant belief-clusters, and the cause of bad epistemic effects is epistemically problematic belief-forming mechanisms. In response, one might reject the claim that echo chambers are neutral. Nguyen, for example, describes echo chambers as "local background conditions that turn generally good epistemic practices into locally unreliable ones", presumably in the sense that they reliably lead to false beliefs on the part of their members (Nguyen 2020: 156). It's not clear, however, that this feature of echo chambers follows from his definition. For example, round-earthers routinely assign low credibility to flat-earthers and high credibility to their critics. Moreover, group membership (belonging to the group critical of flat-earthers) requires a core set of beliefs (representing facts that contradict flat-earther claims) which essentially includes the belief that flat-earthers are unreliable. It would therefore appear that round-earthers occupy an echo chamber according to Nguyen's definition. But since the earth is in fact round, the round-earth echo chamber tends to produce and reinforce true beliefs. This suggests that for any echo chamber built around misinformation, there may be another one built around the truth. This result is at odds with the idea that echo chambers are "local background conditions that turn generally good epistemic practices into locally unreliable ones" (Nguyen 2020: 156) and instead suggests the idea that echo chambers are global background conditions for reinforcing clusters of consonant beliefs whether true or false. This line of reasoning also implies that distrust is not a falsehood-conducive mechanism. It does not have an overall tendency to produce negative epistemic effects but only does so in certain circumstances.

In response, Nguyen could simply add a reliability condition to his definition. Echo chambers would be, roughly put, (a) communities which reinforce distrust through discrediting of outsiders and credential amplifying of insiders which (b) reliably leads

¹¹That's not to say that some of the particular falsehood conducive mechanisms mentioned cannot be part of broader systems that tend to produce true beliefs nor is it to say that we cannot imagine a merely possible world in which such mechanisms *always* produce true beliefs. Moreover, the reason for focusing on the frequency sense rather than a "propensity" sense (Goldman 1979: 11), which describes a mechanism's performance across merely possible worlds, is because the misinformation and echo chamber literature seeks to identify the causal mechanisms which actually cause problems.

¹²By "the cause" I mean the most salient cause in a given context (see Mackie 1965, 1974; Schaffer 2012).

community members to form false beliefs. The problem with this solution is that (b) does all the work in making echo chambers problematic. The solution builds the idea that an echo chamber involves falsehood-conducive mechanisms into the definition, but it doesn't explain how unreliability follows from the formal features of echo chambers described in (a). It also makes the difference between a view which treats echo chambers as neutral and one that treats them as bad seem merely terminological. Putting the point more positively, it suggests the need to distinguish between neutral and negative senses of "echo chamber". This may, however, be a plausible distinction (recall the comparison with the term "conspiracy theory" in section 4).

Nguyen's discussion suggests one other way of capturing the inherent badness of echo chambers without adding reliability conditions: in terms of epistemic vice. At one point, for example, he claims that agents typically enter echo chambers as a result of blameworthy epistemic vice and that echo chambers are typically created and maintained through epistemically malicious intent. The problem with this proposal, however, is that neither phenomenon is a necessary feature of echo chambers even if they are typical (Nguyen 2020: 149). Nguyen would apparently agree with this assessment. In his view, one of the most interesting features of problematic echo chambers is that they can result in bad epistemic effects despite individual members doing their reasonable best to gain knowledge. For this reason, Nguyen claims in another place that echo chambers occur when epistemic vice is a "feature of the collective intelligence, rather than of the individual" (Nguyen 2020: 156). Perhaps this idea could capture the inherent badness of echo chambers, but one problem with this proposal is that many vice epistemologists simply define epistemic vice in terms of reliability (Medina 2013; Cassam 2016). There are alternative conceptions of epistemic vice available (see Battaly 2014, 2015), but Nguyen never clarifies which sense of vice he has in mind and ultimately reverts back to the idea that echo chambers tend to produce bad effects (Nguyen 2020: 156). This suggests a reliabilist interpretation of vice but that would mean his proposal is no different from one that merely appends a reliability condition like (b) to the definition of echo chambers. There may be other proposals for capturing the idea that echo chambers are epistemically bad, but I will continue to use the neutral sense going forward and leave the problem of how to treat them as negative without appending ad hoc reliability conditions as an open challenge.¹³

¹³If Nguyen did make use of one of the alternative conceptions of vice, he may be able to develop a plausible alternative without appealing to ad hoc reliability conditions or construing echo chamber members as necessarily blameworthy for negative epistemic effects. In her treatment of epistemic vice, Heather Battaly distinguishes three notions of epistemic vice: effects vice, personalist vice, and accountability vice (see Battaly 2014, 2015). Roughly put, the first is a reliabilist conception of vice. Something counts as an epistemic vice if and only if it leads to bad epistemic effects. An accountable vice on the other hand is an acquired, blameworthy, and intrinsically bad epistemic property of an agent. Neither fits what Nguyen is after. Personal vices, on the other hand, are person-level and intrinsically bad epistemic properties of an agent, but they are neither acquired through one's own efforts nor are they epistemically blameworthy. As we've seen, putting up credibility filters only leads to bad epistemic effects in particular circumstances, but, as I will explore further in the next section, what's worrisome about problematic echo chambers is that they are highly resistant to correction. By connecting these two lines of reasoning together, perhaps Nguyen could argue that echo chambers are social networks characterized by a vicious form of closed-mindedness, at least (but not exclusively) in a personalist sense. There is a small but recently developing literature on the nature of closed-mindedness, and some authors argue that closed-mindedness is not always epistemically vicious (Fantl 2018; Battaly 2018a, 2018b). Any proposal for making sense of echo chambers in terms of closed-mindedness would therefore have to address that literature. The idea of feedback loops discussed in section 7 would also be relevant. It's still not clear, however, whether Nguyen would accept this proposal because he is attracted to the idea that echo chambers may result in bad effects even when members are trying their best to know the world. Notice this possibility is compatible

7. Collaborative resistance and audio feedback loops

If the cause of bad epistemic effects is the presence of falsehood-conducive mechanisms and echo chambers are more or less universal background conditions for belief-transmission and reinforcement, why are so many of us worried about echo chambers? In this section, I defend the view that echo chambers are worrisome because attempts to correct echo chambers built around false beliefs are met with *collaborative resistance* involving self-reinforcing *feedback loops*. I therefore defend a qualified version of Lackey's claim that echo chambers "are not the problem" (Lackey 2018). Nguyen's treatment again serves as a useful jumping-off point. Nguyen focuses on how echo chambers *create and reinforce* distrust (Nguyen 2020: 146). Expanding on this notion, we can say that echo chambers have the potential to reinforce the very filters which help sustain them in epistemically diverse societies. This makes sense in light of the social network model. The echoes within a social network will include ideas about the topic around which the echo chamber is built but also ideas about who is worth hearing and who is worthy of trust. The repeated expression of these latter commitments will reinforce beliefs about the appropriateness of one's filters. Since filters themselves play a role in belief reinforcement by selectively allowing consonant views through, filters in such cases play a role in reinforcing themselves.

We can make sense of the latter idea using another musical metaphor. Consider the common problem of audio feedback. The sound from a set of speakers is picked up by a microphone, which feeds that sound back through the very same speakers. The speakers then produce a slightly altered version of that sound, which then feeds back into the microphone again, and the process iterates. The feedback loop that results boosts the audio signal and creates an unpleasant squealing sound, which is canceled by either disconnecting the microphone from the speakers or repositioning them so that the speaker-output isn't picked up by the mic. In other words, by breaking the loop. The final metaphor applies to the most interesting and potentially pernicious mechanisms for boosting belief signals in metaphorical echo chambers.

Echo chambers might make use of feedback loops of two different kinds: *intra-network feedback loops* and *inter-network feedback loops*. The first kind occurs between an agent and another member of their echo chamber. Nguyen describes feedback loops of this kind in terms of "bootstrapped corroboration" (Nguyen 2020: 144). With a deficit of dissenting viewpoints and a surplus of corroboration, the agent who deploys robust omission-filters may gain an over-inflated sense of confidence in their view. After all, if you are in a bubble, it appears as though everyone agrees. Bootstrapped corroboration may result in arrogance (perhaps as a domain general trait), but it may simply increase one's conviction in a core belief and provide ancillary beliefs which allow one to express or defend it in new ways.

Bootstrapped corroboration therefore describes a kind of feedback loop which occurs when one seeks out other agents with similar beliefs. Nguyen's description adequately captures the phenomenon, but leaves out one important idea. While the reinforced beliefs might concern what events took place or how some portion of the world works, they might also be about whose voices are worth hearing and who is worthy of trust. The feedback loop mentioned at the beginning of this section is therefore an instance of bootstrapped corroboration which reinforces beliefs about what kind of filters are appropriate.

The second kind of feedback loop occurs between an agent and another agent outside of their echo chamber. According to Nguyen, echo chambers are especially

with my view as an echo chamber may be falsehood-conducive as a whole despite the epistemic virtue of some of its members.

problematic when they involve “disagreement-reinforcement mechanisms” (Nguyen 2020: 147). These are feedback loops which assure that attempts at correction by outsiders only further reinforce one’s core beliefs, including one’s belief that the outsider is untrustworthy (Nguyen 2020: 147–8; see also, Begby 2020). To illustrate this type of feedback loop, Nguyen focuses on conspiracy theories. It’s typically part of a conspiracy theory that the relevant authorities are lying about, for example, who killed JFK or orchestrated the 9/11 attacks. The theory also predicts that the authorities will try to undermine or suppress evidence and discredit dissenting voices. When the authorities inevitably do refute the voice of “truthers”, this is taken as further evidence of a coverup, further indication that the authorities are untrustworthy, and further confirmation of the reliability of conspiratorial voices. Nguyen’s example focuses on the reinforcement of credibility-filters, but we can also imagine inter-network feedback loops reinforcing omission-filters. Sometimes the voices an agent tries to exclude will nevertheless bypass their omission-filters. Given that the agent already disagrees with those voices, further exposure may only reconfirm that they are not worth hearing.

Notice that inter-network feedback loops needn’t rely on echo chambers. The feedback loop occurs when an agent with certain epistemic properties interacts with an agent with contradictory beliefs. The echoing of consonant beliefs throughout that agent’s social network may explain how the first agent originally acquired those epistemic properties, but it doesn’t factor into the explanation of belief reinforcement through the feedback loop. This suggests that only feedback loops of the first kind are central to the topic of echo chambers and that feedback loops of the latter kind are only a typical side-effect. I would also argue that the idea of intra-network feedback loops does a better job of capturing the distinctive danger of echo chambers for two reasons.

First, consider an implication of bootstrapped corroboration. Having certain beliefs leads an individual agent to seek out like-minded agents, and their beliefs are reinforced by giving uptake to the other agent’s belief expressions. By that same token, the other agents within the echo chamber will have their beliefs reinforced by other echo chamber members. This creates the possibility of an interpersonal feedback loop where one’s beliefs are reinforced by being fed through other agents who then feed those same beliefs backwards with increased confidence. When the loop has gone full circle, the agent’s belief has played a non-trivial role in supporting itself by being fed through other agents. As an example, consider the online news and discussion forum Reddit, which allows users to post, upvote, and comment on other users’ posts. The site is a notorious virtual playground for the mechanisms of selective exposure (e.g., you can subscribe to certain subreddits you prefer) and third-party filtering (e.g., the number of upvotes a post receives will determine its prominence on the website) (see, e.g., Marwick and Lewis 2017; Massanari 2017a). A registered member logs onto their preferred subreddit and submits an opinion or a link that expresses their view. Other like-minded users encounter the post, which increases their credence in the linked news story and perhaps boosts their confidence in a general attitude toward some group or idea. Like-minded users then contribute with upvotes, corroborating evidence, similar claims, or expressions of support. These contributions increase the confidence of other users, perhaps including the original poster, who contribute further. This process continues until the thread, for whatever reason, runs out of steam. In sum, the first problem is that multiple agents feed off of each other creating a dynamic and multi-directional form of bootstrapped corroboration.

This leads to the second problem feedback loops create within echo chambers: *collaborative resistance in the face of outsider intervention*. While the inter-personal feedback loops Nguyen describes in terms of “disagreement reinforcement” also play this role, his focus is on the response of an individual agent. Of course, many agents

in an echo chamber may exhibit the sort of response he describes, but in principle the feedback loop only requires one individual inside of an echo chamber and one other agent outside of it in order to take place. Collaborative resistance, on the other hand, is the response of the echo chamber as a collective. Consider the gamergate conspiracy as an example (see, e.g., Kaplan 2014; Parkin 2014). In 2014, Eron Gjoni wrote a blog post detailing a past romantic relationship with independent videogame developer Zoë Quinn. Social media users inferred from Gjoni's blog post that Quinn had slept with a video game journalist in exchange for positive reviews. Although the journalist in question had never reviewed any of Quinn's games, this sparked intense criticism and harassment of Quinn on social media platforms. Eventually the hashtag "#gamergate" was coined on twitter and spread around the internet to express, among other things, the idea that elites within the gaming industry conspired to cover up the purported scandal. While gamergate proponents claimed that the issue was professionalism and ethics in gaming journalism, critics of gamergate viewed it as an excuse for a vitriolic right-wing reaction against the increasing influence of leftists and feminists on gaming culture (O'Rourke 2014; Mantilla 2015; Jane 2016; Massanari 2017a, 2017b).

The momentum of the gamergate campaign quite plausibly relied on interpersonal feedback loops operating across social media platforms which reinforced conspiratorial narratives as well as distrust in the victims and their defenders. Gamergate proponents claimed that the victims of harassment, which eventually included women other than Quinn, were lying, accused the broader gaming industry and mainstream media of a coverup, and helped each other secure purported evidence to back up their claims. Defenders of the victims were frequently labeled "social justice warriors" (SJWs), a derogatory term designed to put their motivations into question, and the community used these labels determine who to trust (Heron *et al.* 2014: 22). While inter-network feedback loops may have played some role in keeping outsiders at bay,¹⁴ the primary mechanism for belief and filter reinforcement was the belief expressions (about what happened and why, about who is worth hearing, and about whom to trust) which reverberated throughout the online echo chamber itself.

In addition to collaborative resistance to outsider intervention, echo chambers also create the problem of *collaborative resistance to insider defection*. While the former describes the critical response of the social network to outsiders, the latter describes the critical response of the social network to the divergent belief-expressions of its members. Internal sanctions carry weight for potential defectors not only because they trust other insiders but also because the various social networks one relies on for epistemic goods are often also relied on for non-epistemic goods. Fellow echo chamber members are our family, friends, co-workers, religious leaders, teachers, and so on. Breaking out of an echo chamber requires putting up various filters against these social networks, at least along epistemic dimensions. While it may be possible to maintain the same non-epistemic relationship with members of one's social group after filtering them out, this isn't always possible. It depends on how central those beliefs are to echo chamber members and their broader social identity. Escaping a deeply problematic echo chamber is therefore not only a matter of securing an alternative source of epistemic goods but also a matter of securing an alternative source of non-epistemic goods (friendship, family, and so on). If no such source is available, one will be unlikely to defect. Echo chamber members will therefore have both epistemic and practical reasons to resist the temptation to defect.

¹⁴For example, calling someone a "SJW" quite plausibly plays a role in "evidential pre-emption" involved in inter-network disagreement reinforcement mechanisms.

How then should one respond to Lackey's charge that echo chambers aren't the problem? In my view, the best response leans on the idea of collaborative resistance. Imagine an echo chamber which lacks substantive truth-conducive mechanisms to control the quality and accuracy of reverberations throughout the social network. Further suppose that members deploy various filters to keep opposing voices at bay. The result would be a situation where any attempted intervention will be met with collaborative resistance and any dissent from within will be met with communal sanctions. This means that interventions aimed at installing corrective mechanisms fail to take hold. On the other hand, a good echo chamber, which already makes use of truth-conducive mechanisms, won't face this problem so long as those mechanisms are well-maintained. Therefore, while the presence or absence of truth-conducive mechanisms separates good and bad echo chambers, a bad echo chamber's resistance to correction is explained in terms of features of the echo chamber itself – particularly the role of collaborative resistance and intra-network feedback loops.¹⁵ If this is right, then at the very least Lackey's claim shouldn't be "echo chambers are not the problem" but that echo chambers aren't the only problem (Lackey 2018).

8. Breaking the loop

If problematic echo chambers are sustained by collaborative resistance and intra-network feedback loops, then successfully addressing them will require breaking the loop. There are two potential targets for intervention: individual echo chamber members or the social network as a whole. In the first case, the goal is to cut an individual off from the problematic social network, at least along epistemic dimensions. Nguyen argues that in particularly extreme cases, escaping the echo chamber requires a *social epistemic reboot* which involves reconsidering "all testimonial sources with presumptive equanimity" and discarding "other background beliefs [arising] from the flawed credential structure of the echo chamber" (Nguyen 2020: 157). Although Nguyen questions whether such an epistemic upheaval is psychologically possible and suggests that it goes beyond the demands of epistemic responsibility, he offers one instructive case. Derek Black was raised from a young age to be a leader in the neo-Nazi community and for a time fulfilled that role. The major impetus for escaping his echo chamber was the persistent and generous friendship of a Jewish classmate in college. The trust between the friends allowed Black to step out of his echo chamber, which led to a systematic reevaluation of his beliefs (compare Ornstein 2016). As Nguyen puts it, "echo chambers work by a manipulation of trust" and so to address them we must "repair the broken trust between echo chamber members and the outside world" (Nguyen 2020: 158–9). This needn't be achieved through a manipulative or coercive form of "deprogramming" but rather by providing repeated and forceful evidence of trustworthiness as a character virtue. This opens the problematic echo chamber member to useful epistemic friction (see Medina 2013), further reexamination of their beliefs, and the development of epistemically useful feedback loops with outsiders.

¹⁵I don't take a definite stand on what the most salient cause of resistance is or whether there really is one. I would speculate that falsehood-conducive mechanisms, collaborative resistance, and filter mechanisms are salient causes in typical cases. This would also mean that, contra Nguyen, we cannot account for the highly recalcitrant nature of problematic echo chambers in terms of distrust alone. My point is not the trivial one that any effect requires multiple causal conditions, but we may not even be able to say that distrust is the most salient cause (see Mackie 1965, 1974; Schaffer 2012). It should also be kept in mind that even though distrust plays a role in preserving false beliefs in a bad echo chamber, this doesn't imply that distrust is a falsehood-conducive mechanism in itself (but closed-mindedness might be, see fn. 13).

As expressed by the particular example, Nguyen's solution is very labor intensive at an interpersonal level and may put too much responsibility on the shoulders of out-group members. Dealing with resistance to insider defection will be particularly challenging. It's hard to maintain a friendship with members of the neo-Nazi community once you've shunned Nazi ideology, and so mounting a rescue attempt of this kind will require securing a new source of non-epistemic goods for the defector, which is no easy task (compare Dennett and LaScola 2015; but see Horton and Freire 1990; Anderson 2013). Alternative interventions may seek to transform the echo chamber from the inside out. Although Benkler *et al.* argue that the "ultimate solution" to the problem of misinformation in contemporary U.S. politics would involve drastic institutional change on the right with a reorientation towards objectivity over partisanship, they offer no comprehensive proposal on how such a "Herculean feat" might be achieved (Benkler *et al.* 2018: 351–5).¹⁶ A step-by-step plan for changing a deeply problematic and highly resistant echo chamber from the inside out may not be possible. In what follows, I limit myself to the general form a solution must take in light of the problem of collaborative resistance.

First of all, because of the problem of collaborative resistance to insider defection, the *first-movers* who initiate change will have to be (or become) highly trusted, visible, and efficacious members of the echo chamber (see Bicchieri and Mercier 2014; Bicchieri 2016; compare O'Connor and Weatherall 2019: 178). In her empirically based treatment of social normative change, Christina Bicchieri argues that social progress requires a simultaneous change in group expectations concerning what other members will do (*empirical expectations*) and what they ought to do (*normative expectations*) and that high-status individuals are particularly effective at instigating change in group expectations (Bicchieri 2016). Candidate first-movers will include news-media institutions, politicians, political pundits, celebrities, influencers, community leaders, collective trendsetters, and so on (Bicchieri 2016: Ch. 5). Because they are high-status members, their deviant behavior will incur fewer or less damaging communal sanctions while generating true-belief echoes throughout the network.¹⁷

Second, because of the problem of collaborative resistance to outsider intervention, first-movers must advocate for change in a way that doesn't undermine their insider status. Bicchieri explores the role of small-scale group deliberation, media campaigns, economic incentives, and legal interventions in promoting changes in shared normative

¹⁶Benkler *et al.* claim that more feasible interventions will target "how journalists pursue their professional commitments in a highly asymmetric media ecosystem" (Benkler *et al.* 2018: 352). They propose modifications to the journalistic norm of balanced reporting and the adoption of new norms governing the coverage of fake news, extremist memes, and email dumps on both the left and right. Although they don't put the point this way, such interventions are designed to dampen the echoes of misinformation in the media-driven echo chamber and to increase exposure to disconfirmation. The dominant "the internet did it" narrative suggests interventions targeting online platforms with the goal of dampening false-belief echoes either by designing algorithms which subtractively filter out misinformation (Pariser 2011) or by adding features which help platform users apply distrust filters correctly (Rini 2017). In *The Misinformation Age*, however, O'Connor and Weatherall warn that fake news stories are not fixed targets and that "whatever barriers we erect against the forces of propaganda will immediately become targets for these sources to overcome" (O'Connor and Weatherall 2019: 175; see also, Benkler *et al.* 2018: Ch. 13).

¹⁷This doesn't mean that high-status agents will be immune to sanctions. If a news media platform, for example, starts to depart from news which is identity-confirming for its base, then it may start to lose that base to others willing to provide such a service. This exact dynamic played out when Breitbart started drawing more attention than Fox News during the 2016 election (Benkler *et al.* 2018: 103, see also Ch. 1, 4, and 5). This problem doesn't undermine the claim that the most likely candidates for initiating change will be high-status agents. But it does emphasize the importance of my second point – that first-movers need to advertise change as in line with the broader values and goals of the echo chamber.

expectations (2016: Ch. 4). In many cases, the interventions that successfully promote normative change target shared normative expectations indirectly. To borrow one of Bicchieri's examples, a soap opera may change social expectations about women's roles within a community by portraying an otherwise prototypical woman who successfully departs from normative expectations. The apparent purpose of the drama is entertainment, but it subtly changes normative expectations by providing a new model for the social role and attendant social scripts. While not as extreme as deprogramming, such interventions count as subtle "nudges" rather than direct reason-giving exchanges (Thaler and Sunstein 2009). To directly advocate for change would be to argue that the need is implied by the group's broader value commitments or that change promotes core goals. Whether communication is direct or indirect, to address the problem of collaborative resistance the change will have to be represented as coming from within the echo chamber.

Bicchieri's account of social normative change focuses on cases where a community is implicitly aware of reasons for change and the threat of internal sanctions is the most salient obstacle to progress. This isn't always the case with problematic echo chambers. While any member will have a reason to change their beliefs simply because their beliefs are false, high-status members may be especially unmotivated to change. The bad epistemic effects of falsehood-conducive mechanisms often fail to produce direct, practical consequences for those most responsible for creating and sustaining them (O'Connor and Weatherall 2019: 176–7), and attempts to help high-status insiders to recognize those consequences from the outside will be met with distrust and collaborative resistance. Part of the problem is therefore *epistemic insouciance* or lack of concern for the truth in favor of practical, especially economic, interests (Frankfurt 2005; Cassam 2018). This suggests that problematic echo chambers will begin to reverberate with true beliefs more often than false beliefs only when the consequences of epistemic carelessness come to disrupt the lives of high-status members.

9. Conclusion

Lackey worries that focusing on echo chambers threatens a "same boat" mentality, but this worry is misplaced. I have argued that echo chambers may be treated as neutral in themselves. While good echo chambers are constrained by truth-conducive mechanisms, bad ones are not. The epistemic relativist or nihilist may reject the distinction between good and bad echo chambers, but the idea that we all occupy echo chambers implies neither view. Leaving these concerns aside, one may still wonder "how do I know I'm in good echo chamber?" To say that one can never know is not to accept relativism or nihilism but a fairly extreme form of skepticism. One that denies, for example, that it is possible for a round-earther to know they are in a good echo chamber. Although I won't attempt a comprehensive anti-skeptical argument here, it may help to briefly address one potential confusion. Given that problematic echo chambers are self-insulating, it's unlikely that one will be able to convince the flat-earther they are wrong. In the face of criticism, they may weaponize the concept of echo chambers to argue that the round-earthers' position is epistemically symmetrical to theirs. But as every teacher knows, there's a difference between knowing and being able to pass knowledge on. The fact that one cannot convince the flat-earther doesn't imply that the two positions are symmetrical. At the same time, the comparison with teaching is instructive. It suggests that epistemic asymmetries shouldn't lead to dismissiveness of the other side or over-confidence in one's own view but to ongoing care in building better echo chambers.

References

- Anderson E.** (2013). *The Imperative of Integration*. Princeton, NJ: Princeton University Press.
- Bakshy E., Messing S. and Adamic L.A.** (2015). 'Exposure to Ideologically Diverse News and Opinion on Facebook'. *Science* **348**(6239), 1130–2.
- Battaly H.** (2014). 'Varieties of Epistemic Vice'. In J. Matheson and R. Vitz (eds), *The Ethics of Belief: Individual and Social*, pp. 51–76. Oxford: Oxford University Press.
- Battaly H.** (2015). 'Epistemic Virtue and Vice: Reliabilism, Responsibilism, and Personalism'. In C. Mi, M. Slote and E. Sosa (eds), *Moral and Intellectual Virtues in Western and Chinese Philosophy*, pp. 99–120. London: Routledge.
- Battaly H.** (2018a). 'Can Closed-mindedness be an Intellectual Virtue?' *Royal Institute of Philosophy Supplements* **84**, 23–45.
- Battaly H.** (2018b). 'Closed-mindedness and Dogmatism.' *Episteme* **15**(3), 261–82.
- Begby E.** (2020). 'Evidential Preemption'. *Philosophy and Phenomenological Research*. Online First. <https://doi.org/10.1111/phpr.12654>.
- Benkler Y., Faris R. and Roberts H.** (2018). *Network Propaganda: Manipulation, Disinformation, and Radicalization in American Politics*. Oxford: Oxford University Press.
- Bicchieri C.** (2016). *Norms in the Wild: How to diagnose, Measure, and Change Social Norms*. Oxford: Oxford University Press.
- Bicchieri C. and Mercier H.** (2014). 'Norms and Beliefs: How Change Occurs'. In M. Xenitidou and B. Edmonds (eds), *The Complexity of Social Norms*, pp. 37–54. New York, NY: Springer.
- Brady W.J., Wills J.A., Jost J.T., Tucker J.A. and Van Bavel J.J.** (2017). 'Emotion Shapes the Diffusion of Moralized Content in Social Networks.' *Proceedings of the National Academy of Sciences USA* **114**(28), 7313–18.
- Cassam Q.** (2016). 'Vice Epistemology.' *The Monist* **99**(2), 159–80.
- Cassam Q.** (2018). 'Epistemic Insouciance.' *Journal of Philosophical Research* **43**, 1–20. Online First. doi: 10.5840/jpr2018828131.
- Chater J.** (2016). 'What the EU Referendum Result Teaches us about the Dangers of the Echo Chamber.' <http://www.newstatesman.com/2016/07/what-eu-referendum-result-teaches-us-about-dangers-echo-chamber>.
- Clark D.** (2018). *Behind the Curve*. Delta-v Productions.
- Craver C. and Tabery J.** (2017). 'Mechanisms in Science.' In E.N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Spring Edition), <<https://plato.stanford.edu/archives/spr2017/entries/science-mechanisms/>>.
- Crockett M.J.** (2017). 'Moral Outrage in the Digital Age.' *Nature Human Behaviour* **1**(11), 769.
- Dennett D.C. and LaScola L.** (2015). *Caught in the Pulpit: Leaving Belief Behind*. Pitchstone Publishing.
- Dotson K.** (2014). 'Conceptualizing Epistemic Oppression.' *Social Epistemology* **28**(2), 115–38.
- Dunning D.** (2011). 'The Dunning-Kruger Effect: On Being Ignorant of One's Own Ignorance.' *Advances in Experimental Social Psychology* **44**, 247.
- Emba C.** (2016). 'Confirmed: Echo chambers Exist on Social Media. So What do we do About Them?' https://www.washingtonpost.com/news/in-theory/wp/2016/07/14/confirmed-echo-chambers-exist-on-social-media-but-what-can-we-do-about-them/?utm_term=.1f5b99e58174.
- Fantl J.** (2018). *Limitations of the Open Mind*. Oxford: Oxford University Press.
- Frankfurt H.G.** (2005). *On Bullshit*. Princeton, NJ: Princeton University Press.
- Fricker M.** (2009). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press.
- Garrett R.K., Weeks B.E. and Neo R.L.** (2016). 'Driving a Wedge Between Evidence and Beliefs: How Online Ideological News Exposure Promotes Political Misperceptions.' *Journal of Computer-Mediated Communication* **21**(5), 331–48.
- Goldman A.I.** (1979). 'What is Justified Belief?' In G. Pappas (ed.), *Justification and Knowledge*, pp. 1–25. Boston: D. Reidel.
- Haugeland J.** (2002). 'Authentic Intentionality.' In M. Scheutz (ed.), *Computationalism: New Directions*, pp. 159–74. Cambridge, MA: MIT Press.
- Heron M.J., Belford P. and Goker A.** (2014). Sexism in the Circuitry: Female Participation in Male-dominated Popular Computer Culture. *ACM SIGCAS Computers and Society* **44**(4), 18–29.
- Hooton C.** (2016). 'Social Media Echo Chambers Gifted Donald Trump the Presidency.' <http://www.independent.co.uk/voices/donald-trump-president-social-media-echo-chamber-hypernormalisation-adam-curtis-protests-blame-a7409481.html>.
- Horton M. and Freire P.** (1990). *We Make the Road by Walking: Conversations on Education and Social Change*. Philadelphia, PA: Temple University Press.

- Hosanagar K.** (2016). 'Blame the Echo Chamber on Facebook. But Blame Yourself, Too.' *Wired*. <https://www.wired.com/2016/11/facebook-echo-chamber/>.
- Jamieson K.H. and Cappella J.N.** (2008). *Echo Chamber: Rush Limbaugh and the Conservative Media Establishment*. Oxford: Oxford University Press.
- Jane E.A.** (2016). *Misogyny Online: A Short (and Brutish) History*. New York, NY: Sage.
- Kaplan S.** (2014). 'With #GamerGate, the Video-game Industry's Growing Pains go Viral.' *The Washington Post*.
- Khalid A.** (2017). 'Tech Creates Our Political Echo Chambers. It Might Also Be A Solution.' <http://www.npr.org/sections/alltechconsidered/2017/04/12/522760479/tech-creates-our-political-echo-chambers-it-might-also-be-a-solution>.
- Lackey J.** (2018). 'True Story: Echo Chambers are Not the Problem.' *Morning Consult*. <https://morningconsult.com/opinions/true-story-echo-chambers-not-problem/>.
- Mackie J.L.** (1965). 'Causes and Conditions.' *American Philosophical Quarterly* 2(4), 245–64.
- Mackie J.L.** (1974). *The Cement of the Universe: A Study of Causation*. Oxford: Clarendon Press.
- Mantilla K.** (2015). *Gendertrolling: How Misogyny Went Viral*. ABC-CLIO.
- Marwick A. and Lewis R.** (2017). *Media Manipulation and Disinformation Online*. New York, NY: Data and Society Research Institute.
- Massanari A.** (2017a). '# Gamergate and The Fapping: How Reddit's Algorithm, Governance, and Culture Support Toxic Technocultures.' *New Media and Society* 19(3), 329–46.
- Massanari A.L.** (2017b). 'Damseling For Dollars': Toxic Technocultures and Geek Masculinity'. In R.A. Lind (ed.), *Race and Gender in Electronic Media: Content, Context, Culture*. New York, NY: Routledge, pp. 316–17.
- McDowell J.** (1996). *Mind and World*. Cambridge, MA: Harvard University Press.
- Medina J.** (2013). *The Epistemology of Resistance: Gender and Racial Oppression, Epistemic Injustice, and Resistant Imaginations*. Oxford: Oxford University Press.
- Moran R.** (2005). 'Getting Told and Being Believed.' *Philosopher's Imprint* 5(5), 1–29.
- Nguyen C.** (2020). 'Echo Chambers and Epistemic Bubbles.' *Episteme* 17(2), 141–61. doi: 10.1017/epi.2018.32.
- O'Connor C. and Weatherall J.O.** (2019). *The Misinformation Age: How False Beliefs Spread*. New Haven, CT: Yale University Press.
- O'Rourke P.** (2014). 'GamerGate has nothing to do with ethics in journalism.' Canada.com.
- Ornstein M.** (2016). *Accidental Courtesy: Daryl Davis, Race and America*. Independent Lens.
- Pariser E.** (2011). *The Filter Bubble: How the New Personalized Web is Changing What we Read and How We Think*. London: Penguin.
- Pariser E.** (2015). 'Did Facebook's Big Study Kill My Filter Bubble Thesis?' *Wired*.
- Parkin S.** (2014). 'Zoe Quinn's Depression Quest.' *The New Yorker*. <https://www.newyorker.com/tech/annals-of-technology/zoe-quinn-depression-quest>.
- Rini R.** (2017). 'Fake News and Partisan Epistemology.' *Kennedy Institute of Ethics Journal* 27(2), E-43.
- Schaffer J.** (2012). 'Causal Contextualisms.' In M. Blaauw (ed.), *Contrastivism in Philosophy: New Perspectives*, pp. 43–71. London: Routledge.
- Spinney L.** (2017). 'The Shared Past that Wasn't.' *Nature News* 543(7644), 168–70.
- Sunstein C.R.** (2001). *Republic.com*. Princeton, NJ: Princeton University Press.
- Sunstein C.R.** (2009a). *Going to Extremes: How Like Minds Unite and Divide*. Oxford: Oxford University Press.
- Sunstein C.R.** (2009b). *Republic.com 2.0*. Princeton, NJ: Princeton University Press.
- Thaler R.H. and Sunstein C.R.** (2009). *Nudge: Improving Decisions about Health, Wealth, and Happiness*. London: Penguin.
- Tufekci Z.** (2014). 'Engineering the Public: Big Data, Surveillance and Computational Politics.' *First Monday* 19(7). doi: 10.5210/FM.V19i7.4901.

Benjamin Elzinga is an instructor at Georgetown University. His research centers around epistemology, philosophy of mind, and free will, and he is especially interested in the relationship between know how and learning.