

Responding to the Spread of Conspiracy Theories

Nader Shoaibi

September 11, 2022

Abstract

Conspiracy theories are spreading faster than ever and pose a real danger to our societies. It is natural to accuse the consumers of conspiracy theories of irrationality – that they are either not looking at or appropriately sensitive to all the available evidence. In this paper, I attempt to determine if we can make sense of this general idea. I argue that we cannot: conspiracy theories do not spread because the people who believe them are irrational – at least, not necessarily so. In addition, I explore some alternative strategies for responding to the problem of the spread of conspiracy theories. I argue that in addition to confrontational strategies such as social shaming, we need more constructive programs of community activism to battle the spread of conspiracy theories.

Keywords - *conspiracy theories, echo chambers, shaming, activism, social epistemology, disagreement*

1 Introduction

Almost immediately after swarms of rioters stormed the Capitol building on Jan 6, 2021, there were rumors that the whole ordeal was a ‘false flag’ operation by Antifa meant to sully the name of Trump supporters and undermine their accusations of widespread voter fraud.¹ Amusing, if you ask some. Unfortunately for all, we have learned the hard way that conspiracy theories such as this are not to be taken lightly. Consider the now-infamous events of the Pizzagate which led to shots being fired over children’s heads as they dined with their families in their local pizza parlor.² When it comes to some conspiracy theories, the stakes could not be higher. What is more, the innocuous-looking conspiracy theories such as those concerning popular culture (e.g., ‘Elvis faked his death’) or those that

¹Anderson (2021). As I edit this manuscript, the events of the Uvalde shooting are shaking the nation to the core. However, this hasn’t stopped similar disappointingly predictable ‘false flag’ theories to spread on the internet suggesting that the shooting was staged by powerful forces ‘out of a playbook’ with malicious political motives (Siders 2022).

²See Clary, Allman, and Bohn (2016).

seem outright silly (e.g., ‘birds aren’t real’) are also problematic. For, studies have shown that the best predictor of who believes in a conspiracy theory is belief in other conspiracy theories.³ It seems like believing any conspiracy theory – silly or not – is a risk factor in believing other conspiracy theories. Simply ignoring conspiracy theories, then, is not an option because they either pose imminent harm or, short of that, prime the believer to believe other potentially harmful conspiracy theories.

So, what *should* we do about conspiracy theories? Some might argue that the solution is easy: make sure that people do not fall prey to fallacies and biases of critical thinking, that they look at all the relevant evidence and properly adjust their beliefs accordingly.⁴ However, this assumes that the belief in conspiracy theories involves failures of rationality in handling the available evidence.

In this paper, my first aim is to determine if this assumption stands up to scrutiny. I argue that ultimately it does not: sometimes people end up believing in conspiracy theories due to no (rational) fault of their own, and that they cannot avoid this problem by deploying better thinking habits or critical thinking. The more positive aim of the paper is to explore some alternative responses to the problem of the spread of conspiracy theories. In particular, I aim to defend a response that meets an attractive basic desideratum: that it not contribute to the toxic hyper-polarization that we witness in our communities these days.⁵

Here is the plan: I begin in §2 by tackling the question of what conspiracy theories are. Drawing on the recent literature on the concept of ‘fake news’, I identify what I take to be an especially important species of conspiracy theories. My claim, in a nutshell, is that conspiracy theories – or one species thereof anyway – is fake news which involves a conspiracy. In §3, I shift the discussion from what conspiracy theories are to how they spread. I consider two models: *epistemic bubbles* (§3.1) and *echo chambers* (§3.2). While the former attempts to explain the spread in terms of lack of exposure to relevant evidence, the latter does so in terms of the distrust of outside voices. Both models can appear to offer what we might call *rational* diagnoses of the problem: diagnoses that locate the problem with a rational failure to handle one’s evidence. However, I argue that in both cases the diagnoses fail to do full justice to the problem of the spread of conspiracy theories, and that therefore the problem of the spread of conspiracy theories is not necessarily a failure of rationality. Our discussion of the echo chamber model, in addition, reveals an important feature of the spread of conspiracy theories: i.e., that they rely on and exploit people’s *trust*. In §4, I utilize this insight to explore some alternative strategies of responding to the spread of conspiracy theories. While many have found a strategy of socially shaming the consumers of conspiracy theories attractive, I defend a hybrid view which combines social shaming with community activism.

³See Swami et al. (2011). Cited in Cassam (2019).

⁴See, for instance, Cassam (2016).

⁵See Sanders (2021).

2 What are Conspiracy Theories?

The first thing to notice about conspiracy theories is that the words ‘conspiracy’ and ‘theory’ do not do justice to what we mean when we call something a conspiracy theory. To see this, consider the official story of 9/11: a small number of extremist individuals worked in secret to launch an attack on multiple sites in the United States. This is a theory about the events of 9/11 that involves a conspiracy.⁶ We might even go as far as calling it a theory about a conspiracy. But hardly anyone would call it a *conspiracy theory*. To call it that would be to imply that there is something defective about the story.⁷ Exactly what this defect is we are trying to pinpoint by calling some story a conspiracy theory is something we would do well to get clear on, but the point for right now is that we all recognize that the phrase ‘conspiracy theory’ is pejorative: it indicates some sort of defect in the story. To avoid confusing conspiracy theories in the pejorative sense from any old theory about a conspiracy, let us capitalize the first letters of the former and refer to them as ‘Conspiracy Theories’ (CTs for short) from here on out.⁸

Now, what is distinctive about CTs then? Let me start by saying that I don’t see much hope for a traditional analysis of the concept of CTs. Accordingly, in what follows, I don’t take on the ambitious task of providing a series of necessary and jointly sufficient conditions for CTs. What I do aim to provide is the demarcation of one important species of CTs, one that, I think, is particularly salient in our political climate. For everything that I will say, there may be other forms of CTs that deserve our attention, but I will limit myself to the species that I demarcate here.

There is a growing contingency of philosophers who have offered analyses of the notion of ‘fake news’, and I think we can shed light on an important species of CTs by looking at this literature⁹ While I want to remain neutral on the details here, it is helpful to establish some common ground about the notion of fake news. I believe there’s something of a consensus that can be captured by the following three conditions:

1. Doxastic Condition: fake news must not be believed by the creator to be true
2. Intention Condition: fake news must be published with the aim of deceiving a large audience
3. Publication Condition: fake news must be published in traditional media outlets

It is worth noting that as I’ve formulated it the Doxastic Condition — arguably the most central of all the conditions — is consistent with two different kinds of analyses: those according to which the creator of fake news believes the story to be false (as defended by Rini

⁶For our purposes, we can understand a conspiracy to be a plot by a small group of individuals who are working together in secret to do something harmful.

⁷Some philosophers, notably, Dentith (2014), defend a revision of the concept of conspiracy theory such that it sheds this negative connotation. For the purposes of this paper, I ignore such a view. For critical discussion, see Napolitano (2021).

⁸I am following Cassam (2019) here.

⁹See, for instance, Blake-Turner (2020), Lackey (2021), Mukerji (2018), and Rini (2017) among many others.

(2017)), on the one hand, and those according to which the creator is simply indifferent to the truth or falsity of the claims she makes, which makes her a Frankfurtian bullshitter (as defended by Mukerji (2018)).¹⁰ My proposal is simply to appropriate this analysis for CTs:

CTs are conspiracy stories that are not believed by their creators to be true and are propagated with the aim of deceiving a large audience.

One modification that seems necessary is that unlike fake news, since by their very nature conspiracies are secretive and often involve government officials, stories that purport to reveal CTs need not be published in traditional media. Rather, CTs have a way of going around in a more organic way through, for instance, the word of mouth and face-to-face interactions. So, we have reason to drop the publication condition and only appropriate the doxastic and intention conditions for our analysis of CTs.

It might be argued that this definition of CTs exclude many clear cases. For instance, consider the CT that the moon landing was fake. Plausibly, many propagators of this CT believe it to be true and don't have the intention of misleading anyone. In response, first, it is worth emphasizing that here I am drawing a crucial distinction between the *creators* of CTs and their *consumers*. Many consumers of CTs end up believing them and, partly for that reason, are able to propagate them even more effectively. However, according to this definition, that is not a relevant factor in determining whether a story is a CT. What is important is that the original creators of the story fabricate the CT, either believing it to be false or being indifferent to its truth or falsity.

Still, this might leave some unsatisfied because arguably there are CTs whose original creators did believe them to be true. Wouldn't those kinds of cases constitute a counterexample to our definition? This brings me to the second, and more important, point about this definition, which is that the aim of the definition is not to offer a traditional analysis of the concept of CT but to demarcate one important species of it. While it may be true that some CTs start innocently in the sense that their creators genuinely believe them to be true, the suggestion is that it remains an interesting fact that for many others the creators don't believe them to be true, and that is sufficient to demarcate a special species of CTs for us.

To appreciate the interest in this kind of CT, consider the CT that the Sandy Hook shooting was fake. Alex Jones started this CT only hours after the first reports were coming out. During his live show on Infowars on December 14, 2012, as he covered the story, he began airing the idea that the shooting was staged, and that the shooter was part of a government mind control program.¹¹ Jones' accusations on that day are a clear case of conspiracy mongering that blatantly dispenses with evidence or argument. As he later confessed, he did not believe these theories to be true. I submit that this is not an isolated kind of CT. In our present political climate, those CTs that are pieces of fake news (in that their creators don't believe them to be true and propagate them with the intention of creating deception)

¹⁰For Frankfurt's original account of bullshit see Frankfurt (2005).

¹¹See Kirk et al. (2020).

constitute an important species of CTs in general which deserve our attention.¹²

Before I move on to our diagnosis of the problem of the spread of CTs, let me briefly address an immediate worry. It might be argued that our analysis reduces CTs to fake news, and that therefore there is little reason to pursue CTs as a separate phenomenon. While much of what I have to say in the rest of the paper applies to fake news just as well, I think that there is a good reason to specifically focus on CTs. As we will see in the next section (§3.2), the root of the problem of the spread of both fake news and CTs is the misalignment that they manage to create and maintain in terms of whom the consumers of social media tend to trust. There are, as we shall see, different mechanisms that can be utilized to create this misalignment. However, CTs have an explicit mechanism to do this: namely, *automatic epistemic preemption*. This makes CTs ideal for diagnosing the problem of the spread of misinformation and reflecting on viable ways of responding to it.

3 The Problem of the Spread of Conspiracy Theories

According to our definition, CTs are defective theories that involve a conspiracy story because their creators either believe them to be false or are indifferent as to whether they are true. In short, they are defective because they are pieces of fake news. Does this mean that the *consumers* are always in the wrong to believe such CTs? In the next two subsections, I examine the hypothesis that there is something rationally amiss in the consumers of CTs. I offer two models of the spread of CTs, and explore some potential diagnoses of what epistemic shortcomings might lead one to fall for them. I argue that none of these diagnoses is entirely satisfactory: some fail to identify a genuine shortcoming, while others can't explain the full extent of the problem of the spread of CTs. The upshot will be that consumers of CTs are not necessarily epistemically irrational for believing CTs.

3.1 The Epistemic Bubbles Model

Here is one attractive explanation of the spread of CTs. We live at a time when social media and messaging platforms such as TikTok, Twitter, WhatsApp, Telegram, etc. have become staples of everyday life. These platforms, however, utilize selection mechanism that create an environment in which people end up believing CTs. There are roughly three related issues here: First, people on social media tend to trust their 'friends' and connections, and therefore end up readily accepting their posts and comments; second, because one tends to see the same story shared by multiple sources within one's social network, this creates a sense of confidence that the story must be true; and, third, within one's social media certain voices tend to be left out which creates a vacuum of contrary evidence. In all these cases, the spread of CTs is explained by way of an information network from which certain

¹²As I mentioned in footnote 1, in the wake of the Uvalde shooting, there appears to be no shortage of baseless claims and CTs about the incident. Arguably, many of these constitute similar cases where the creators don't believe them to be true. Alex Jones, for instance, only 2 days after the incident, has suggested that the shooting was orchestrated by 'opportunistic motives' (Kaonga 2022).

voices are omitted. Let's call such explanations the *epistemic bubble model* of the problem of the spread of CTs.¹³

It is worth noting that we can distinguish two kinds of selection mechanisms that contribute to an environment from which certain voices are absent: On the one hand, there are the external factors having to do with the design and implementation of social media and messaging platforms. For instance, filtering and hyper-customization on these platforms attempt to deliver exactly the kinds of stories and posts that the user finds interesting or desirable. On the other hand, there are agent-centered factors such as self-selection, confirmation bias, etc. (Nguyen 2020, p. 144).

The epistemic bubble model of the spread of CTs suggests three distinct epistemic diagnoses of the problem:

1. People tend to trust the 'friends' and connections on social media, but they should know better.
2. People tend to take repeated exposure to a story on their social media as further evidence in its support, but they should know better.
3. People tend to ignore the lack of contrary evidence on social media, but they should know better.

Not all these diagnoses, however, succeed. Let us examine each in turn.

Grant, first, that people tend to accept the posts and comments of their 'friends' on social media. Is there something epistemically problematic about this? To be sure, it is epistemically wrong to blindly give credibility to those who we have deemed to be our epistemic peers. For instance, if we are considering the question of how much a certain bag of fruits weighs, it would be a mistake to take one person's word over another based on whether they share our political viewpoint. However, as Rini (2017) points out, there are many domains in which such a partiality for our epistemic peers is epistemically acceptable. For instance, arguably in perceptual domains what allows us to rely on others' testimony is the assumption that their perceptual systems are sufficiently like ours, or, to take another example, in normative domains, we can reasonably use another person's partisan affiliations as proxy for whether they are to be trusted or not.¹⁴ The question for our purpose is whether social media platforms constitute such a domain. I agree with Rini that the answer is 'yes' because, like in normative domains, it is reasonable to use another's position on important social and political issues as proxy for whether they are to be trusted on similar issues.¹⁵ When I see my reformist (in the context of the Iranian politics) friend's post

¹³My usage of epistemic bubbles roughly follows Nguyen (2020), who distinguishes between epistemic bubbles and echo chambers. Unlike the former phenomenon in which certain voices are left out, in echo chambers certain sources are discredited. As we shall see in the next subsection (§3.2), the latter constitutes a much more interesting phenomenon and calls for a more involved treatment, which we will turn to below.

¹⁴See Rini (2017, E51-2) for relevant citations for these examples.

¹⁵Worsnip (2019, pp. 7–8) claims that normative reliability of some source does not support its descriptive reliability, and so normative disagreement is not grounds for trusting a source on its factual assertions. However, he fails to appreciate that normative reliability is what we might call a 'general' feature of someone's character: it shows that the individual has goodwill and is in the business of doing the right thing.

on Facebook about the newly released CIA documents revealing that Ayatollah Kashani was not only opposed to the duly elected Prime Minister Mosadegh but was in close communication with the US and even requested money from them,¹⁶ I assume reasonably that she has legitimate sources for her post and that therefore I have reason to accept the claim. The fact that people tend to trust their connections on social media, then, serves only as a reminder that people tend to, entirely within their epistemic rights, trust each other based on proxy information about their normative viewpoints. So, trusting one's connections on social media in and of itself is not irrational.

Let us turn now to the second of the three diagnoses that the epistemic bubble explanation offers: that people tend to get further confirmation of their views based on what they see on their social media platforms. The problem, according to this diagnosis, is that it is an epistemic mistake to allow the sheer number of endorsements to influence one's confidence in a situation in which the endorsements are not independent of each other.¹⁷ At first, this might seem reasonable enough. However, exactly how we understand the notion of *dependence* here matters. For instance, as Lackey (2021) has recently argued, to the extent that dependent believers endorse a claim autonomously – that is, through a reflective process of engaging with the evidence and its source which gives rise to epistemic responsibility – then the numbers should indeed increase our confidence. After all, every time the claim survives the scrutiny that each endorsement involves that provides further evidence in support of the claim (210-211).¹⁸

Admittedly, however, we cannot reasonably assume that many social media consumers are exactly autonomous in the above sense. The real trouble for this diagnosis of the problem of the spread of CTs, again as argued convincingly by Lackey (*ibid.*), lies in the fact that even if people fall short of autonomous dependence, they may nevertheless be excellent at telling whether a given source is trustworthy or not. So, even if a given agent matches her views exactly with what she hears from some source, that still gives her testimony the ability to provide further evidence for those views because they have survived her powers of discrimination with respect to the source (213). So, given that, as we discussed above, it is reasonable for me to trust my connections and their discriminatory powers with respect to which sources to trust, when they share and reshare the same story, that gives me reason to increase my confidence in the story. Their testimony has the power to add evidential value to the story that they share.

This, then, brings us to the third diagnosis that the epistemic bubble explanation offers: that there is a lack of reliable coverage of the relevant evidence. Note, first, that being selective about one's sources is not bad in and of itself. There is plenty of noise out there, the wholesale uptake of which would make it impossible for any finite being to do much

This is different from reliability about matters of fact, which is domain-specific. If, for instance, someone is a reliable mechanic, that shows nothing one way or another about their reliability about other domains. However, if someone has a good character, then that does seem to constitute good reason to trust them in matters of thought and knowledge.

¹⁶Allen-Ebrahimian 2021.

¹⁷This point is familiar from the debate on epistemology of disagreement. See, for instance, (Kelly 2010).

¹⁸The debate goes back to (Goldman 2001), which famously argued that non-independent testimony should be discounted. For other challenges against this view, see (Lackey 2013).

else.¹⁹

However, the problem is that there are certain *relevant* arguments, observations, and ideas that are likely left out. For instance, if my social media is saturated with CTs about, say, the involvement of Antifa in the Jan 6 storming of the Capitol, I would likely completely miss out on the coverage of the rioters' trials and the evidence that is brought before the courts. That is relevant contrary information that I'm simply deprived of in virtue of inhabiting a certain social network. Given this situation, I would be making an epistemic mistake if I were to form the belief that there is no such evidence based on what I am seeing on social media.

It is true that often we can reasonably respond to the demand for ruling out the existence of contrary evidence by 'But I would have seen evidence to the contrary if there were any such evidence'. However, given the fact that social media tends to leave out relevant voices, one is not epistemically justified in making that assumption in the context of social media sharing.

And the problem is not just about social media. As Worsnip (2019) has recently argued, we have reason to believe that there are 'illicit' influences on all news media outlets about what they report and how much importance they give them. Worsnip identifies what he calls 'inconvenient truths', that is, truths that are in conflict with the news media outlets' standing beliefs and commitments. According to him, inconvenient truths create incentive for the outlets to either report falsely on or simply ignore those inconvenient truths. Thus, Worsnip argues, 'reading only news sources from any one side of the political spectrum (even if that side is normatively correct) will result in getting a skewed, incomplete picture of the evidence' (15). If Worsnip is right, then, it's not just on social media that we have reason to expect a lack of diverse voices on a given topic. News media corporations, too, often fail to expose us to the full array of relevant information.

I believe that while the first two diagnoses of the problem of the spread of CTs that epistemic bubble explanation offers fail, with the third we have identified a genuine problem: consumers of CTs lack coverage-reliable beliefs because epistemic bubbles filter out relevant evidence.

Now, given this diagnosis of the problem, the solution seems obvious. Since, as we have seen, the source of the problem is essentially lack of information access, the proper response will have to simply focus on making sure that all the relevant information is available on our social media, massaging platforms, and one's choice of news outlets.²⁰ To be clear, when I say that the solution is easy, I don't mean to suggest that it would not require a lot of effort, expertise, and ingenuity: I have no doubt that there are very difficult problems about design and implementation of mechanisms that are meant to address the information access problem. And while we may be already doing some things to address the issues, the problem is far from being solved.²¹ The solution is easy only in the sense that *theoretically* we know exactly what needs to be done. I say that this is suspicious be-

¹⁹See Lackey (2021) for more on the issue of noise.

²⁰On the news media side, Worsnip (2019) argues that we have an obligation to diversify our sources.

²¹For a recent development, see Wagner (2021).

cause given the attention that this problem has attracted lately, one would expect there to be not much disagreement about this. However, there are aspects of the spread of CTs on social media and messaging platforms, which we have not even begun to touch on. In the rest of this section, I turn to an alternative diagnosis of the problem of the spread of CTs, which aims to reveal some of these aspects.

3.2 The Echo Chamber Model

Nguyen (2020) has recently distinguished epistemic bubbles from ‘echo chambers’. According to Nguyen, while, as we have seen, epistemic bubbles involve leaving out certain voices as a result of different selection mechanisms, the latter is a much more robust and pernicious phenomenon in which the views from outside the echo chamber are systematically discredited and the members’ trust is instead directed toward select narratives. Echo chambers, thus, work like cults: They recruit members by earning their trust and maintain their dominance by reinforcing the insider narrative while undermining other sources from the outside.²²

Compared to epistemic bubbles, echo chambers offer a much more accurate explanation of the spread of CTs (call it the *echo chamber model*). That is because they account for the reluctance of CTs to change their views in the face of contrary evidence. This is a widely noted feature of CTs. For instance, Cassam (2019), contends that one of the six defining features of CTs is that they are ‘self-sealing’:²³ i.e., they are ‘difficult for outsiders to challenge’.²⁴ Similarly, Napolitano (2021) proposes to define CTs as beliefs in conspiracies that are ‘self-insulated’ in the sense that the belief is ‘immune to being disconfirmed by the kind of evidence that is available in normal circumstances’ (88). Echo chambers nicely account for this feature of conspiracy theories because by their very nature they target the members’ trust and seek to insulate the insider beliefs from outside influence. So, if conspiracy theories do spread primarily through echo chambers, then it is no surprise that they would end up being self-sealing in the sense that their members’ tend to be reluctant to change their views based on contrary evidence.

It can seem as if echo chambers also provide a ready diagnosis of the problem of the spread of CTs. One might argue, for instance, that since one must fit one’s beliefs to the available evidence, the Conspiracy Theorist’s reluctance to adjust her beliefs to contrary evidence is irrational and therefore epistemically wrong.²⁵ The idea is that being epistemically *recalcitrant* cannot be rational when one possesses explicit evidence that speaks strongly enough against one’s beliefs. CTs seem to be particularly prone to this kind of irrationality because they are notoriously adamant in their conspiracy beliefs and are always ready to find evermore intricate ways of dismissing any evidence that goes against their beliefs.

It is worth emphasizing that this diagnosis implies that the problem with the spread of CTs

²²See Nguyen (2018) for the analogy between echo chambers and cults.

²³The other five features are: *speculative, contrarian, esoteric, amateurish, and pre-modern*.

²⁴Cassam in turn is following Sunstein and Vermeule (2009) in identifying CTs with the self-sealing feature.

²⁵For a recent example of this argument, see Napolitano (2021).

is not exhausted by the lack — or unavailability — of contrary evidence. Consumers of CTs are often diligent evidence gatherers, who seek those who disagree with their views. They engage the other side and actively tweak their arguments. The root of the problem, according to this diagnosis, is rather that consumers of CTs fail to adjust their beliefs in light of the contrary evidence that admittedly they are often careful to gather. This, according to the diagnosis, is a failure of rationality.

As convincing as this diagnosis may seem, it fails to identify a general problem with the spread of CTs. To see this, let us examine echo chambers more closely. One of the interesting features of echo chambers is that they can create rational pathways for their members to maintain their inside beliefs even in the face of explicit contrary evidence. Consider, for instance, the phenomenon that Begby (2021) has identified as ‘epistemic preemption by neutralization’: the process by which the speaker of a testimony preempts the evidence by signaling to the hearer ahead of time that the evidence has already been taken into account. For Begby, this signaling takes the form of an explicit warning: ‘My opponents will tell you that q , but I say p ’, where p and q are in more-or-less obvious tension with each other. His example is Bill O’Reilly, warning his audience that President Obama might appear to undermine his analysis of falling wages by saying that it is due to the recession. As Begby notes, the aim here is to somehow undercut the President’s opposing view, and his analysis is that epistemic preemption by neutralization is what allows O’Reilly to do this.

Begby convincingly argues that epistemic preemption by neutralization can be rational because of the following justification transmission principle: ‘whatever justifies my epistemic trust in my testifier’s assertions that p also justifies my epistemic trust in my testifier’s assessment of the total evidence that bears on p ’ (519). The idea is that given the transmission principle, the testifier can exploit her favorable epistemic position by indicating to the hearer that subsequent apparently contrary evidence is already taken into account. When this happens, the evidence that would have had the effect of cancelling one’s reason to believe that p had it been presented to one without forewarning, is ‘rendered epistemically moot’ (*ibid*).

Epistemic preemption by neutralization constitutes one mechanism by which the trust dynamics of an echo chamber can create a resistance to contrary evidence.²⁶ In particular, it relies on the considerable trust that the members of the echo chamber put into certain voices inside the echo chamber. When O’Reilly, for instance, warns his audience that President Obama will try to convince them that falling wages are due to recession, and that in fact the issue is the result of machines replacing people, he is exploiting the trust that he knows his base has in him.²⁷ It is this trust that allows him to preempt the future evidence that his audience will get in the form of testimony from President Obama because, given the trust, his audience is justified in believing that he has already considered the relevant evidence.

²⁶It is worth emphasizing that it would be a mistake to understand the language of ‘resistance’ and ‘recalcitrance’ psychologically; it is crucial that epistemic preemption by neutralization and its kin are mechanism that create a *rational* resistance to evidence in the sense that under the non-ideal circumstances that we all find ourselves, one is not failing, epistemically speaking, to be as one should be.

²⁷See “Bill O’Reilly” (2015). Cited in Begby (2021).

This is already bad news for the diagnosis that the echo chamber explanation of the spread of CTs makes. However, the situation is worse. For, as Begby argues, there is an even more powerful variation of the mechanism of epistemic preemption that allow the members of echo chambers to maintain their inside beliefs in the face of explicit contrary evidence. This mechanism relies not only on the trust that the members of the echo chamber have in certain inside voices, but also on the *distrust* that is encouraged of outside sources. Here is how it works: The trusted inside voice warns the members that outside sources will attempt to undermine the members' belief by providing contrary evidence and arguing that the inside voices are not trustworthy. As we have seen, one effect of this kind of warning is to neutralize the subsequent contrary evidence. However, when the outside voices do attempt to provide contrary evidence, this also provides the members with evidence that the inside voices are indeed trustworthy simply because their predictions come to be true. So, in what has been labeled as an act of 'intellectual judo',²⁸ echo chambers have the power to turn the very presentation of contrary evidence from discredited outside sources into further reason to hold the original beliefs. Following Begby, Let us call this *epistemic preemption by discrediting*.

As far as I can tell, the first mention of something resembling epistemic preemption (not under this label) as a means of the spread of CTs goes back to Keeley (1999):

By invoking a conspiracy hypothesis, large amounts of 'evidence' are thrown into question. This is one of the most curious features of these theories: to my knowledge, conspiracy theories [i.e., explanations involving conspiracies] are the only theories for which *evidence against them is actually construed as evidence in favor of them*. The more evidence piled up by the authorities in favor of a given theory, the more the conspiracy theorist points to how badly 'They' must want us to believe the official story. (120, *emphasis is mine*)

Others have also noted this 'curious' feature. For instance, citing studies on the so-called 'backfire effect', Cassam (2019, ch. 4) argues that CTs are immune to purely intellectual interventions because sometimes corrections and contrary evidence have the effect of *strengthening* the original beliefs.²⁹

Now, one might interpret claims like these as pointing merely to (defective) psychological mechanisms that enable belief in CTs. However, the point I want to make, following Begby, is a normative one: that in the right context, (1) consumers of CTs can legitimately dismiss contrary evidence and (2) the interventions to undermine CTs provide further evidence for the CT. Let's see how exactly that is supposed to work.

Take the example of Alex Jones and his CT that the Sandy Hook Shooting was a 'false flag' operation by powerful elements in the government meant to provide justification to

²⁸See Nguyen (2018).

²⁹While a proper treatment of Napolitano (2021) goes beyond the scope of this paper, it is worth noting that even though she takes note of the mechanism of preemption by discrediting, she confesses that 'I find it hard to make sense of Keeley's claim that seemingly disconfirming evidence could be construed as supporting evidence, rather than just as neutral evidence' (fn. 20). I believe that this reveals a detrimental blindness in her overall argument, but I cannot argue the point here.

take away American guns. Alex Jones' theory, as any other CT, involves a conspiracy: that certain individuals (perhaps even agencies) within the US government plot together in secret to do something harmful. Because the conspirators in a conspiracy operate in secret, there is good reason to expect them to attempt to provide seeming evidence against there being a conspiracy. This basic idea plays a crucial role for Conspiracy Theorist like Alex Jones. To see this, consider someone for whom the appropriate epistemic conditions have been satisfied such that they are justified in trusting Alex Jones.³⁰ For this individual, given their trust in Alex Jones' assertions, the subsequent evidence from government sources is moot. After all, that information is thought to be coming from the conspirators, and as such it is already signaled as misleading or otherwise accounted for.

What is more, given that the conspirators by definition work in secret, the very stipulation of a conspiracy by Conspiracy Theorists such as Jones implies that there will likely be attempts to cover up the purported conspiracy. When the government sources deny the existence of a conspiracy in presenting their account of the events, this provides consumers with further evidence for the CT and the trustworthiness of its creator. In presenting a CT, Conspiracy Theorists like Jones effectively make a prediction: the conspirators will attempt to cover up their plots. When this prediction comes 'true' in the form of government officials denying the existence of a conspiracy, it constitutes further confirmation that the inside voices are trustworthy.

What we have with the case of CTs, then, is a new modality of epistemic preemption, one which operates *automatically*, without the need for the Conspiracy Theorist to make any *explicit* predictions about what the supposed conspirators would say or do. While, as we have seen, epistemic preemption – by neutralization and by discrediting – is a powerful tool for the spread of any information through echo chambers, the new modality of epistemic preemption — let us call it *automatic epistemic preemption* — is even a more powerful tool that specifically can help to spread CTs. Any CT has two immediate implications: (1) since the conspirators work in secret, any future evidence can be explained away as something from an untrustworthy source or something already accounted for (we can call this *automatic preemption by neutralization*), and (2) since the conspirators would attempt to undermine the CT, any future attempt to undermine the CT is taken as further evidence in support of the CT (we can call this *automatic preemption by discrediting*).

All of this, however, raises a crucial question: Can anyone really be justified in putting their trust in the inside voices of a CT? For instance, is there really any room left, given the overwhelming evidence,³¹ to say that someone in their right mind might be justified in trusting Alex Jones? I agree that when it comes to people like Alex Jones, the answer is likely 'No'. It is scarcely my intention here to offer an excuse for Conspiracy Theorists like Alex Jones and their followers. However, it is important to realize that under not so farfetched scenarios, people can be entirely within their epistemic rights to trust the inside voices of a CT. In the literature on *epistemological permissivism*, i.e., the view that for some body of evidence there is more than one rational doxastic attitude any agent with that

³⁰For the record, while I highly doubt that such an individual actually exists, the point remains that someone in the right context could be justified in trusting Jones. I will return to this issue below.

³¹See Kirk et al. (2020).

evidence can take, it is often argued that people's upbringing, for instance, can put them in epistemological positions such that they can have radically different ways of assessing and weighing the evidence.³² This gives us a recipe for cases in which trusting the inside voices of a CT seems entirely rational: Consider someone who is born into a CT. From early childhood, they are exposed to the underlying assumptions of this CT – for instance, they learn the identities of the conspirators and begin to develop an eye for detecting their plots and designs. It should be clear that this kind of education can create an epistemic environment for the individual such that they non-culpably rest their trust in the inside voices of a CT and thereby become susceptible to mechanisms of epistemic preemption.

Let us take stock. Over the last two subsections, we considered two models of the spread of CTs, and examined some potential diagnoses of the rational shortcomings that could lead one to believing a CT. However, we found that none of these diagnoses is entirely satisfactory.

I think it is fair to conclude that the problem of the spread of CTs is *not* a problem of epistemic rationality — not at least necessarily so.³³ Believing a conspiracy theory, in other words, is not something that requires some kind of epistemic shortcoming. One can be entirely within their epistemic rights and still end up believing a conspiracy theory. Of course, many do end up believing conspiracy theories because they fail to consider or properly take in the relevant evidence. And for that, they may well be guilty of epistemic irrationality. However, crucially, this need not be the case. People, for instance, who have been, out of no fault of their own, inculcated into a CT from early childhood and groomed to put their trust in the inside voices can end up believing CTs without any kind of epistemic shortcoming. As far as the problem of the spread of CTs goes, we are not always facing a problem of epistemic irrationality.³⁴

Despite this conclusion about the *spread* of CTs, it is important to remember that CTs are problematic for multiple reasons. As we have seen, creators of CTs intend to deceive their audience through creating theories that they do not believe to be true. Deception is bad for our society³⁵ – not least because it can lead to dangerous behavior such as those we witnessed in the Pizzagate case. Furthermore, given that CTs can combine with morally corrupt and nefarious motives, they pose even more serious dangers to our societies. The

³²See, for instance, Schoenfield (2014) among many others. For criticism, see, for instance, Dogramaci and Horowitz (2016).

³³For a relatively early endorsement of a similar conclusion, see Baurmann (2007). Baurmann's term for echo chambers is 'fundamentalist equilibrium' and he argues that groups who are in such epistemic networks may end up believing fundamentalist ideologies without any violations of rationality.

³⁴One might object that even the groomed Conspiracy Theorist is falling short of an epistemic ideal: They believe something to be true for which there is good contrary evidence. However, that would constitute a change of subject. The relevant notion of rationality for our purposes is one whose primary role is in cases of *second*-person advice. This notion of rationality minimally requires that it be deployable in giving constructive advice. The notion of rationality that this objection relies on, however, is merely *classificatory* that categorizes beliefs as good or bad, and thus departs from our intended notion. For more discussion of the distinction between different senses of normativity see, for instance, Kolodny (2005) and Thomson (2008) among others.

³⁵It is beyond the scope of this paper to expand further on this point. See Blake-Turner (2020) and Lackey (2021) among others for further discussion in the context of discussing fake news.

challenge, then, is this: Given that the consumers of CTs are not necessarily epistemically irrational, and thus the response strategy cannot be to encourage ways of avoiding irrationality through, for instance, further critical thinking or mitigating cognitive biases, what might a viable response strategy look like? This is the question I turn to in the next section.

4 Responding to the Spread of Conspiracy Theories

Our discussion of the echo chamber model of the spread of CTs reveals that an important feature of CTs is the *trust structure* that they create for their consumers: CT consumers *trust* select inside voices at the cost of *discrediting* all outside sources of evidence. It is this trust structure that allows for CTs to take advantage of the mechanisms of epistemic preemption and to thereby insulate the consumers from contrary evidence. This points to a general strategy of responding to CTs: *incentivize the consumers to reorient their trust*. The idea is that if, on the one hand, the consumers reassess their trust in the inside voices, and, on the other, establish trust in other sources of evidence, this can interrupt the cycle of epistemic preemption and the spread of CTs.

In the context of discussing fake news, Blake-Turner (2020) suggests one way of achieving this. He claims that we can shame people away from believing and propagating fake news by socially exposing and sanctioning them:³⁶

Reproaching one another for our individual failings is a powerful tool. To stop the proliferation of fake news, and all the ensuing epistemic problems, we need to start taking one another to task for engaging with fake news in ways that degrade our epistemic environment. ... Social sanction is an effective method of norm enforcement and internalization (13).

We can propose the same strategy against the consumers of CTs: One powerful way of mitigating the spread of CTs is through calling each other out for believing in CTs, with the purpose of creating a feeling of shame. We can call this the *social shaming strategy*. Of course, given our conclusion in the previous section, i.e., that there is no failure of rationality in consumers of CTs, this strategy has the burden of explaining what justifies socially exposing consumers of CT.

Cassam (2019) also defends the social shaming strategy. He argues that the proper response to CTs will involve 'outing' consumers of CTs for, on the one hand, the ultimate function of their theories, which, according to him, is *political propaganda*, and, on the other, the morally corrupt views that underpin many CTs.³⁷ The latter is specifically crucial in Cassam's view because it is the mechanism by which we can create a feeling of shame in

³⁶Another proponent of this kind of strategy is Lackey (2021).

³⁷It might be argued that while both *outing* and *shaming* involve socially exposing an individual for the view that they hold or embody, only shaming is done with the intention of creating a feeling of shame in the individual. Strictly speaking, on this way of thinking, exposing someone without an accompanied intention to create a feeling of shame (even if it incidentally leads to it) is not a case of shaming. As it should be clear

those who believe CTs but also those who might otherwise flirt with them. For instance, he thinks that the anti-Semitic viewpoint from which many CTs originate is something that must be pointed out and made explicit in order to shame people away from CTs.

This is where what I have described as the ‘outing’ of Conspiracy Theories as anti-Semitic or as right-wing propaganda becomes important. The point of trying to show that many Conspiracy Theories are anti-Semitic or peddling right-wing propaganda is to shame or embarrass people into not flirting with them. Shame and embarrassment are *emotions*. If a proto-Conspiracy Theorist is embarrassed to discover that Conspiracy Theories have historical links with anti-Semitism, her embarrassment is an expression of her *values*. The truth is that values count for more than abstract reasoning in political debate. If anti-Semitism is the original sin of Conspiracy Theories, that is something that good people will care about. Not caring about it is a sign that one is *morally* deficient (74).

It is worth emphasizing that this kind of strategy can only work if the target individuals share roughly the same moral outlook, according to which the CTs involve morally corrupt views. If a Conspiracy Theorist does not share the moral outlook according to which antisemitism is not morally wrong, then exposing them for holding such views will not succeed in creating a feeling of shame.

In addition, the following line of thought suggests that the social shaming strategy shares the justificatory burden of Blake-Turner’s proposal: The same kind of grooming that can allow an individual to rationally believe in a CT can also allow an individual to hold *morally* reprehensible views without culpability. If this is right, then Cassam similarly owes us an explanation of what justifies calling out the consumers of CT.

Drawing on Calhoun (1989), Blake-Turner offers an interesting response to the justificatory burden. The response begins with distinguishing between *normal* and *abnormal* moral contexts. The distinction is meant to track widespread recognition and compliance with correct moral norms. Those contexts in which there is widespread recognition and compliance are *normal*, whereas those in which there is not are *abnormal*. The next move is to argue that in abnormal contexts, it is hard to excuse non-culpable violators without thereby condoning their behavior. Excusing a scientist, for instance, for his immoral behavior based on the widespread practice of engaging in that behavior is tantamount to justifying it. The upshot is that in abnormal contexts there may be license to call out violators of correct moral norms, even if they are not culpable because they could not have known better. Since we are in an abnormal moral context regarding CTs, that is, we are in a situation in which people explicitly endorse, for instance, the anti-Semitic views as part of their belief in CTs, we have license to call consumers of CT out even if they are not culpable for the moral views that they endorse.

from my discussion so far, I’m focusing on the concept of shaming here, and, as the quote below shows, though Cassam sometimes uses the label ‘outing’, his proposed strategy is also shaming, that is, exposing someone with the intention of creating a feeling of shame.

Blake-Turner appropriates this response for explaining what justifies calling out believers in fake news based on their epistemic costs, and we can extend that to the case of consumers of CTs. Thus, we can distinguish between normal and abnormal *epistemic* contexts, where the former indicates widespread recognition and compliance with epistemic norms and the latter indicates lack thereof. Since, as we have noted, CTs involve misleading claims that are not believed by their creators to be true, this will situate us in an abnormal epistemic context. But then, given that abnormal contexts allow for exposing violators of correct norms — even if they are not culpable — we have justification for socially shaming the consumers of CTs.

I think that this response elegantly addresses the burden of justifying shaming non-culpable consumers of CT. However, it leaves an important underlying issue with this strategy intact. In the introduction (§1), I said that a minimal desideratum of a response to the problem of the spread of CTs is that it doesn't contribute to the toxic division that we are witnessing in our societies these days. And therein lies the fundamental issue with social shaming as a strategy of responding to CTs. Social shaming may be a power tool to get people to stop openly talking about CTs, but it can also easily deepen the mistrust in outside sources because they appear hostile and antagonistic. If that is right, then social shaming is likely to lead to *worsening* the problem of the spread of CTs. For, in light of the *echo chamber model* of the spread of CTs, it should be clear that deepening of the mistrust in outside sources is exactly what is behind the spread of CTs. Social shaming, then, at least in the bare-bones version that I have presented it here, cannot be a viable strategy to respond to the spread of CTs.

An alternative strategy of incentivizing the reorientation of trust in the consumers of CTs is what we can label 'community activism'. Allow me to introduce this strategy with an example that I find particularly compelling. In the first months after the COVID vaccines were made available in 2020, there was a striking level of vaccination hesitancy in the minoritized communities and in particular the Black American community.³⁸ However, while the political divide in vaccinations persists,³⁹ the racial gap has largely subsided.⁴⁰ There are two questions about this situation that are pertinent for our purposes: (1) what was the root cause of the problem of vaccination hesitancy in the Black American community?, and (2) what helped to mitigate the problem?

A recent New York Times article suggests revealing answers to both questions.⁴¹ According to the article, one of the underlying causes of the hesitancy was 'a powerful combination of general mistrust of the government and medical institutions, and misinformation over the safety and efficacy of the vaccines'. It documents the experiences of individuals from Tuskegee, AL, where a 40-year study by a government medical institution has left the community sacred and frustrated. The 'Tuskegee Study of Untreated Syphilis in the Negro Male' as it was called by the U.S. Public Health Service began in the early 1930s and aimed at documenting the natural history of syphilis. The participants were offered

³⁸See Walker et al. (2021).

³⁹See Kirzinger et al. (2021).

⁴⁰See Ndugga et al. (2021).

⁴¹See Burch and Walker (2021).

'free medical exams, free meals, and burial insurance' in exchange for participating in the study. By the 1940s, Penicillin became the standard treatment for syphilis, winning the Nobel Prize in Medicine and Physiology in 1945 for its discoverers, Alexander Fleming, Ernst Boris Chain and Sir Howard Walter Florey. The infuriating fact is, however, that the treatment was held from the participants in the Tuskegee study – not until 1972 when an investigation by the Assistant Secretary for Health and Scientific Affairs finally brought it to an end.⁴²

Using our terminology, we can say that horrible experiences such as the Tuskegee experiment help to create echo chambers, i.e., networks of mistrust in outside sources of information, in which CTs can take root and spread. But what helped to mitigate the problem of vaccination racial gap? The same article offers the following:

How the racial gap was narrowed — after months of disappointing turnout and limited access — is a testament to decisions made in many states to send familiar faces to knock on doors and dispel myths about the vaccines' effectiveness, provide internet access to make appointments and offer transportation to vaccine sites.

I believe that this can provide us with a model for an alternative answer to the spread of CTs in general. What we need is a community of individuals working together to reach out to reiterate, argue, and engage the consumers of CTs such that the trust structures that contribute to the spread of CTs are recreated anew. An important aspect of this bottom-up approach that bears emphasizing is the individual-level attention and engagement. One of the reasons that people end up joining echo chambers and similar network such as cults is a lack of social and communal sense of belonging. It is the often desperate attempt to fill this gap that leads many to join such networks. The strategy of community activism attempts to target exactly this lack. And it is this feature of the strategy that sets it apart from the social shaming strategy: while the latter tends to deepen a sense of isolation from the community, community activism tries to overcome it.

In his treatment of echo chambers, Nguyen (2020) makes a similar suggestion. He cites the story of Derek Black, who was groomed by his father to be a neo-Nazi leader but was able to 'escape' his social network and publicly renounced white nationalism. In his recounting of what led to his transformation, Black emphasizes his encounter with several Jewish people, in particular, Matthew Stevenson, who invited Black into his Shabbat dinners and initiated a years-long conversation with him. Black's description of this encounter fits the 'community activism' model: he was given a face-to-face opportunity to understand the perspective and ideas from a source outside his restricted social network. According to Black, Stevenson's aim was not to try to convince him, but just to 'let me see a Jewish community thing so that if I was going to keep saying these anti-Semitic things that at least I had seen real Jews'.⁴³

A striking feature of Black's story, however, is that it actually combines the two strategies

⁴²See *Tuskegee Study - Timeline - CDC - NCHHSTP* (2021).

⁴³Barbaro 2017.

that we have been looking at in this section. While Black emphasizes the community activism aspect of his story through his encounter with Stevenson, he is also clear that his transformation first began with something else: ‘It was after I got outed on campus. First, this response of condemnation.’⁴⁴ Understandably, given the explicit racist nature of his former beliefs, the condemnation took a moral flavor: that his world view is corrupt and that he does not belong to the campus community. Black’s story, then, is an example of the social shaming and community activism strategies working together.

This hybrid position may well be the position we should prefer. Given the trust structure that the spread of CTs relies on, the response to the spread of CTs must have two aspects: one that addresses the *trust* in the inside voices and one which addresses the *distrust* of outside sources. The strategy of social shaming seems best suited for the first aspect because it forces the consumer of CT to dislodge her trust in the inside voices. However, as we noted above, this strategy on its own is likely to deepen the epistemic gap between the consumer of CTs and the outside sources. What is needed is an attempt to establish new connections and welcome the individual into the larger society. This is where community activism comes into play. While social shaming interrupts the normal epistemic economy for the consumer of CTs, community activism creates a new network of epistemic connection that allow the individual to take in the relevant evidence and form new beliefs accordingly.

5 Conclusion

The problem of the spread of CTs is not necessarily a problem of epistemic rationality – or so I have argued in this paper. I have considered multiple attempts to argue that there is something irrational about the consumers of CTs believing their theories. On the one hand, while the *epistemic bubble model* does correctly identify a problem of lack of relevant information in some social networks, it nevertheless fails as a general diagnosis of the problem because it does not capture a crucial feature of CTs, namely, that the consumers of CTs are recalcitrant in the face of contrary evidence (§3.1). The *echo chamber model*, on the other hand, fits better with our understanding of CTs’ different features; however, it still fails to provide a rational diagnosis of the problem of the spread of CTs because the mechanism of *automatic epistemic preemption* make it possible for consumers of CTs to rationally remain recalcitrant in the face of contrary evidence (§3.2).

If the problem of the spread of CTs was a problem of rationality (or lack thereof), then at least on a theoretical level the solution would be easy: we would just have to fix those rational shortcomings in one way or another. However, if, as I have tried to show, CTs don’t necessarily spread because of failures of rationality, then that seems to leave us with no strategy to respond to CTs. In the last section (§4), I tried to show that in fact the insight from the echo chamber model can be helpful in identifying some strategies for responding to CTs. I explored two suggestions. I argued that the most obvious strategy, i.e., social shaming, by itself is not viable. However, a hybrid strategy that combines social sham-

⁴⁴Hunter-Gault 2019.

ing with a program of community activism seems to be a promising way of incentivizing consumers of CTs to reorient their trust and transform out of CTs.

References

- Allen-Ebrahimian, Bethany (2021). *64 Years Later, CIA Finally Releases Details of Iranian Coup*. Foreign Policy. URL: <https://foreignpolicy.com/2017/06/20/64-years-later-cia-finally-releases-details-of-iranian-coup-iran-tehran-oil/> (visited on 12/26/2021) (cit. on p. 7).
- Anderson, Meg (Mar. 2, 2021). “Antifa Didn’t Storm The Capitol. Just Ask The Rioters.” In: NPR. *Investigations*. URL: <https://www.npr.org/2021/03/02/972564176/antifa-didnt-storm-the-capitol-just-ask-the-rioters> (visited on 01/05/2022) (cit. on p. 1).
- Barbaro, Michael (Aug. 22, 2017). “‘The Daily’: A Conversation With a Former White Nationalist”. In: *The New York Times. Podcasts*. URL: <https://www.nytimes.com/2017/08/22/podcasts/the-daily/former-white-nationalist-derek-black.html> (visited on 01/05/2022) (cit. on p. 17).
- Baurmann, Michael (June 2007). “Rational Fundamentalism? An Explanatory Model of Fundamentalist Beliefs1”. In: *Episteme* 4.2, pp. 150–166 (cit. on p. 13).
- Begby, Endre (2021). “Evidential Preemption”. In: *Philosophy and Phenomenological Research* 102.3, pp. 515–530 (cit. on p. 10).
- “Bill O’Reilly” (May 26, 2015). “Bill O’Reilly: The World According to President Obama”. In: *Fox News*. URL: <https://www.foxnews.com/transcript/bill-oreilly-the-world-according-to-president-obama> (visited on 01/02/2022) (cit. on p. 10).
- Blake-Turner, Christopher (Feb. 10, 2020). “Fake News, Relevant Alternatives, and the Degradation of Our Epistemic Environment”. In: *Inquiry* 0.0, pp. 1–21 (cit. on pp. 3, 13, 14).
- Burch, Audra D. S. and Amy Schoenfeld Walker (Oct. 13, 2021). “Why Many Black Americans Changed Their Minds About Covid Shots”. In: *The New York Times. U.S.* URL: <https://www.nytimes.com/2021/10/13/us/black-americans-vaccine-tuskegee.html> (visited on 01/05/2022) (cit. on p. 16).
- Calhoun, Cheshire (1989). “Responsibility and Reproach”. In: *Ethics* 99.2, pp. 389–406 (cit. on p. 15).
- Cassam, Quassim (Apr. 1, 2016). “Vice Epistemology”. In: *The Monist* 99.2 (2), pp. 159–180 (cit. on p. 2).
- (2019). *Conspiracy Theories*. Think. Cambridge, UK ; Medford, MA: Polity Press. 127 pp. (cit. on pp. 2, 3, 9, 11, 14).
- Clary, Kevin, Greg Allman, and Daniel Bohn (Dec. 5, 2016). “Gun-Brandishing Man Sought to Investigate Fake News Story Site, Police Say — CNN Politics”. In: *CNN*. URL: <https://www.cnn.com/2016/12/04/politics/gun-incident-fake-news/index.html> (visited on 12/30/2021) (cit. on p. 1).
- Dentith, M. (Nov. 6, 2014). *The Philosophy of Conspiracy Theories*. Springer. 203 pp. (cit. on p. 3).
- Dogramaci, Sinan and Sophie Horowitz (2016). “An Argument for Uniqueness About Evidential Support”. In: *Philosophical Issues* 26.1, pp. 130–147 (cit. on p. 13).
- Frankfurt, Harry G. (2005). *On Bullshit*. Princeton, NJ: Princeton University Press. 67 pp. (cit. on p. 4).

- Goldman, Alvin I. (2001). "Experts: Which Ones Should You Trust?" In: *Philosophy and Phenomenological Research* 63.1, pp. 85–110 (cit. on p. 7).
- Hunter-Gault, Charlayne (Nov. 5, 2019). *Derek Black Grew up as a White Nationalist. Here's How He Changed His Mind*. In collab. with Derek Black. URL: <https://www.pbs.org/newshour/show/derek-black-grew-up-as-a-white-nationalist-heres-how-he-changed-his-mind> (visited on 01/05/2022) (cit. on p. 18).
- Kaonga, Gerarrd (May 25, 2022). "Alex Jones Questions Texas School Shooting: "Very Opportunistic"". In: *Newsweek*. URL: <https://www.newsweek.com/alex-jones-texas-school-shooting-robb-elementary-sandy-hook-joe-biden-1710033> (visited on 05/29/2022) (cit. on p. 5).
- Keeley, Brian L. (1999). "Of Conspiracy Theories". In: *The Journal of Philosophy* 96.3, pp. 109–126 (cit. on p. 11).
- Kelly, Thomas (Aug. 5, 2010). "Peer Disagreement and Higher-Order Evidence". In: *Disagreement*. Ed. by Richard Feldman and Ted A. Warfield. Oxford University Press, pp. 111–174 (cit. on p. 7).
- Kirk, Michael et al., director (2020). *United States of Conspiracy*. URL: <https://www.pbs.org/wgbh/frontline/film/united-states-of-conspiracy/> (visited on 01/03/2022) (cit. on pp. 4, 12).
- Kirzinger, Ashley et al. (Nov. 16, 2021). "KFF COVID-19 Vaccine Monitor: The Increasing Importance of Partisanship in Predicting COVID-19 Vaccination Status". In: *KFF*. URL: <https://www.kff.org/coronavirus-covid-19/poll-finding/importance-of-partisanship-predicting-vaccination-status/> (visited on 01/05/2022) (cit. on p. 16).
- Kolodny, Niko (2005). "Why Be Rational?" In: *Mind* 114.455, pp. 509–563 (cit. on p. 13).
- Lackey, Jennifer (2013). "Disagreement and Belief Dependence: Why Numbers Matter". In: *The Epistemology of Disagreement: New Essays*. Ed. by David Phiroze Christensen and Jennifer Lackey. Oxford, UK: Oxford University Press, pp. 243–268 (cit. on p. 7).
- (2021). "Echo Chambers, Fake News, and Social Epistemology". In: *The Epistemology of Fake News*. Ed. by Sven Bernecker, Amy K Flowerree, and Thomas Grundmann. New York: Oxford University Press, pp. 206–227. URL: <https://dx.doi.org/10.1093/oso/9780198863977.001.0001> (visited on 11/26/2021) (cit. on pp. 3, 7, 8, 13, 14).
- Mukerji, Nikil (Dec. 11, 2018). "What Is Fake News?" In: *Ergo, an Open Access Journal of Philosophy* 5.20201214 (cit. on pp. 3, 4).
- Napolitano, Giulia (2021). "Conspiracy Theories and Evidential Self-Insulation". In: *The Epistemology of Fake News*. Ed. by Sven Bernecker, Amy K Flowerree, and Thomas Grundmann. New York: Oxford University Press, pp. 82–105. URL: <https://dx.doi.org/10.1093/oso/9780198863977.001.0001> (visited on 11/26/2021) (cit. on pp. 3, 9, 11).
- Ndugga, Nambi et al. (Dec. 15, 2021). "Latest Data on COVID-19 Vaccinations by Race/Ethnicity". In: *KFF*. URL: <https://www.kff.org/coronavirus-covid-19/issue-brief/latest-data-on-covid-19-vaccinations-by-race-ethnicity/> (visited on 01/05/2022) (cit. on p. 16).
- Nguyen, C. Thi (2018). *Escape the Eco Chamber*. Aeon. URL: <https://aeon.co/essays/why-its-as-hard-to-escape-an-echo-chamber-as-it-is-to-flee-a-cult> (visited on 11/05/2021) (cit. on pp. 9, 11).

- Nguyen, C. Thi (June 2020). "Echo Chambers and Epistemic Bubbles". In: *Episteme* 17.2, pp. 141–161 (cit. on pp. 6, 9, 17).
- Rini, Regina (2017). "Fake News and Partisan Epistemology". In: *Kennedy Institute of Ethics Journal* 27.2, E–43 (cit. on pp. 3, 6).
- Sanders, Linley (Jan. 20, 2021). "Americans Now See Other Americans as the Biggest Threat to Their Way of Life — YouGov". In: *YouGovAmerica*. URL: <https://today.yougov.com/topics/politics/articles-reports/2021/01/20/americans-now-see-other-americans-biggest-threat> (visited on 01/06/2022) (cit. on p. 2).
- Schoenfeld, Miriam (June 2014). "Permission to Believe: Why Permissivism Is True and What It Tells Us About Irrelevant Influences on Belief: Permission to Believe". In: *Noûs* 48.2, pp. 193–218 (cit. on p. 13).
- Siders, David (May 27, 2022). "'It's Straight out of a Playbook': At NRA Convention, Conspiracy Theories Abound". In: *POLITICO*. URL: <https://www.politico.com/news/2022/05/27/nra-convention-uvalde-shooting-00035842> (visited on 05/29/2022) (cit. on p. 1).
- Sunstein, Cass R. and Adrian Vermeule (June 2009). "Conspiracy Theories: Causes and Cures*". In: *Journal of Political Philosophy* 17.2, pp. 202–227 (cit. on p. 9).
- Swami, Viren et al. (Aug. 2011). "Conspiracist Ideation in Britain and Austria: Evidence of a Monological Belief System and Associations between Individual Psychological Differences and Real-World and Fictitious Conspiracy Theories". In: *British Journal of Psychology* (London, England: 1953) 102.3, pp. 443–463 (cit. on p. 2).
- Thomson, Judith Jarvis (2008). *Normativity*. Open Court Publishing. 240 pp. (cit. on p. 13).
- Tuskegee Study - Timeline - CDC - NCHHSTP* (May 3, 2021). CDC. URL: <https://www.cdc.gov/tuskegee/timeline.htm> (visited on 01/06/2022) (cit. on p. 17).
- Wagner, Kurt (Aug. 17, 2021). "Twitter Will Let Users Report Misinformation for First Time". In: *Bloomberg.com*. URL: <https://www.bloomberg.com/news/articles/2021-08-17/twitter-will-let-users-report-misinformation-for-first-time> (visited on 01/06/2022) (cit. on p. 8).
- Walker, Amy Schoenfeld et al. (May 14, 2021). "The Racial Gap in U.S. Vaccinations Is Shrinking, but Work Remains". In: *The New York Times*. U.S. URL: <https://www.nytimes.com/interactive/2021/05/14/us/vaccine-race-gap.html> (visited on 01/05/2022) (cit. on p. 16).
- Worsnip, Alex (2019). "The Obligation to Diversify One's Sources: Against Epistemic Partisanship in the Consumption of News Media". In: *Media Ethics: Free Speech and the Requirements of Democracy*. Ed. by Carl Fox and Joe Saunders. London: Routledge, pp. 240–264 (cit. on pp. 6, 8).