



Inquiry

An Interdisciplinary Journal of Philosophy

ISSN: 0020-174X (Print) 1502-3923 (Online) Journal homepage: <https://www.tandfonline.com/loi/sinq20>

Fake news, relevant alternatives, and the degradation of our epistemic environment

Christopher Blake-Turner

To cite this article: Christopher Blake-Turner (2020): Fake news, relevant alternatives, and the degradation of our epistemic environment, *Inquiry*, DOI: [10.1080/0020174X.2020.1725623](https://doi.org/10.1080/0020174X.2020.1725623)

To link to this article: <https://doi.org/10.1080/0020174X.2020.1725623>



Published online: 10 Feb 2020.



Submit your article to this journal [↗](#)



View related articles [↗](#)



View Crossmark data [↗](#)



Fake news, relevant alternatives, and the degradation of our epistemic environment

Christopher Blake-Turner 

Department of Philosophy, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA

ABSTRACT

This paper contributes to the growing literature in social epistemology of diagnosing the epistemically problematic features of fake news. I identify two novel problems: the problem of relevant alternatives; and the problem of the degradation of the epistemic environment. The former arises among individual epistemic transactions. By making salient, and thereby relevant, alternatives to knowledge claims, fake news stories threaten knowledge. The problem of the degradation of the epistemic environment arises at the level of entire epistemic communities. I introduce the notion of an epistemic environment, roughly the totality of resources and circumstances relevant to assessing the epistemically interesting statuses, such as knowledge. Fake news degrades our epistemic environment by undermining confidence in epistemic institutions and altering epistemic habits, thereby making the environment less conducive to achieving positive epistemic statuses. This is problematic even if the decrease in confidence and the altering of habits are rational. I end by considering solutions to these problems, stressing the importance of reproaching each other for proliferating fake news. I argue that we should reproach even faultless purveyors of fake news. This is because fake news typically arises in abnormal epistemic contexts, where there is widespread ignorance of, and noncompliance with, correct epistemic norms.

ARTICLE HISTORY Received 30 December 2018; Accepted 18 July 2019

KEYWORDS Social epistemology; fake news; relevant alternatives; epistemic environment; reproach

Fake news has gained increasing attention from epistemologists in the last few years, and for good reason. As writers such as Levy (2017) and Rini (2017) point out, fake news is epistemically problematic in various respects. This paper contributes to this body of literature, by highlighting and analysing two novel and epistemically threatening features of fake news. I begin by characterising fake news in §1. I then draw on the relevant alternatives framework in §2 to argue that fake news is problematic even if it isn't believed. A piece of fake news can make relevant an alternative that,

CONTACT Christopher Blake-Turner  chrisbt@live.unc.edu  Philosophy Department, University of North Carolina at Chapel Hill, Caldwell Hall, 240 East Cameron Ave, Chapel Hill, NC 27599, USA

© 2020 Informa UK Limited, trading as Taylor & Francis Group

if the agent is not in a position to rule out, makes it harder for her to gain knowledge. This is *the problem of relevant alternatives*.

Next, in §3, I introduce the notion of an epistemic environment: the circumstances, resources, and other factors of an epistemic community that determine whether one of its members is in a position to gain positive epistemic statuses. This notion brings to light *the problem of the degradation of the epistemic environment*. Fake news makes the conditions of our epistemic community less conducive to achieving knowledge, understanding, rational belief, and so on.

Finally, I sketch some suggestions for what to do about all of this in §4. Although the upshot is ultimately somewhat pessimistic, I stress the importance of criticising purveyors of fake news, even when they are not culpable. This recruits social condemnation as a powerful tool of norm enforcement and internalisation.

1. Fake news

Philosophers have proposed several plausible definitions of fake news (Rini 2017; Aikin and Talisse 2018; Gelfert 2018; Mukerji 2018; Pepp, Michaelson, and Sterken, [forthcoming](#)). My argument goes through on any of these definitions; all that's required is that fake news provides agents with alternatives to claims they might know. Nonetheless, it will be helpful to ground the discussion that follows by focusing on one characterisation. To that end, I follow Rini's (2017, 45) definition of fake news:

FAKE NEWS. A fake news story is one that purports to describe events in the real world, typically by mimicking the conventions of traditional media reportage, yet is [not justifiably believed by its creators to be significantly true], and is transmitted [by them] with the two goals of being widely re-transmitted and of deceiving at least some of its audience.

Rini's definition has several advantages. It requires that there be some dissemblance in the production of fake news, thereby correctly excluding erroneous but good-faith reporting. Moreover, by locating the deception with the creators of the story, it allows fake news to be spread in good faith as long as it originates with some deception.

FAKE NEWS departs from Rini's definition in one important respect.¹ The original formulation requires the creators of a fake news story to know that

¹At least from its letter. Rini (2017, n. 6) herself flags the point about bullshit that I'm about to raise.

the story is significantly false – that is, false in most respects. I require only the weaker condition that producers of fake news *not justifiably believe that the story is significantly true*. There are two reasons to weaken Rini's definition in this way. First, fake news can be originated not only as an outright lie, but also as bullshit in Frankfurt's (2005) sense. Bullshit is characterised by an indifference to truth, a much lower epistemic bar than knowledge of falsehood. Second, the amended definition allows agents to be on the hook for fake news that arises due to epistemic recklessness, even if it happens to be true. Such recklessness might fall short of bullshitting: it embodies insufficient regard for the truth, rather than outright indifference.

2. Relevant alternatives

2.1. The framework

The relevant alternatives framework is a way of capturing what it is for an agent to know something.² Its core claim is the following.

(RA) *S* knows that *p* if and only if *S* is in a position to rule out all the relevant alternatives to *p*.

The intuitive case for (RA) is this. I know something just in case there aren't competing hypotheses. I don't know that the butler did it if my evidence leaves it open that the valet did it. But suppose the only live options are the butler, the valet, and the gardener. Then if my evidence eliminates both the valet and the gardener, then I know that the butler did it.

Several features of (RA) are worthy of remark. First, I intend it to be read as a mere biconditional, compatible with various explanatory stories concerning its key concepts. For instance, (RA) can be read as a piece of conceptual analysis, understanding knowledge in terms of the concepts on the right hand side. While this is a natural understanding of (RA), it need not so be read. Hence it is compatible with a knowledge-first approach

²Several classic formulations of the view are clear that the focus is as much on the theory of knowledge, as on the theory of knowledge ascription (Dretske 1970; Stine 1976; Cohen 1988; Lewis 1996). Hence, as I shall understand it here, the relevant alternatives framework is concerned primarily with whether a subject knows that *p*, rather than whether it is correct to ascribe knowledge to her. So it is not in itself a form of semantic contextualism (see DeRose 1992, 1995; Hawthorne 2003; Nagel 2008). The theory of knowledge is also what I'm interested in here, though given the T-schema equivalence between "*S* knows that *p*" is true' and '*S* knows that *p*', there is of course important interaction between theories of knowledge and theories of knowledge ascription. The taxonomic contours – for instance between invariantism, contextualism, and relevant alternatives theories – are complicated (Stoutenburg 2017). Rather than getting bogged down in classificatory subtleties, I intend my characterisation of the relevant alternatives view to be stipulative.

to epistemology, according to which knowledge is not ultimately analysable in more primitive terms (Williamson 2000). A partisan of that approach could acknowledge the truth of (RA), while denying that the notion of *ruling out a relevant alternative to p* is characterisable independently of *knowing that p*.

Second, I follow Lewis (1996) in giving a modal construal of the concept of an alternative: *p* and *q* are alternatives just in case it is not possible that both *p* and *q*. Thus *I am exactly five feet tall* and *I am exactly six feet tall* are alternatives.

Third, to be in a position to rule out an alternative is to have evidence that is incompatible with it. Because nothing much turns on it here, I remain neutral on the notion of evidence. In particular, I allow that one can have evidence that *p* without being in a position to know that one has evidence that *p*: evidence need not be luminous in Williamson's (2000) sense.

Finally, I need to say something about what makes an alternative *relevant* or not. (RA) does not claim that knowing that *p* requires eliminating every alternative to it. Let's revisit the case of the butler, the valet, and the gardener to bring this out. By supposing that those three are the only live possibilities, I stipulate that the relevant alternatives to *the butler did it* are *the valet did it* and *the gardener did it*. But one might object: surely it's *possible* that someone else did it, which is all that the modal construal of an alternative requires. For instance, perhaps the butler has a twin sister who framed her. Once this possibility is made salient, the evidence that rules out the valet and the gardener (witness testimony, say) does not rule out that the twin did it (the witness can't discriminate between the twins). So, without further evidence, my claim to knowing that the butler did it seems to collapse under this pressure.

Two features of this toy example are worth emphasising. First, whether an alternative is relevant or not depends on the non-epistemic context.³ By 'non-epistemic context', I mean features of the subject's context not traditionally appealed to in analyses of knowledge.⁴ In particular, making an alternative salient suffices to make it relevant (Lewis 1996, 559–560). By mentioning the butler's twin, I draw attention to an alternative that was not previously relevant. In this new context, my evidence is insufficient for my knowledge claim: given that there is now a relevant alternative

³For a different kind of view that might lay claim to the title of a 'relevant alternatives' view (but see n. 2) that is not context dependent in this way see Rysiew (2001, 2007).

⁴Since I am taking the relevant alternatives framework to be a theory of knowledge, rather than a form of semantic contextualism, I focus on the subject's context, rather than the ascriber's. See n. 2.

(the twin did it) that I'm not in a position to rule out, I do not know that the butler did it.

Second, the toy case has the structure of a more traditional skeptical challenge. An initial claim to knowledge is made: *the butler did it; I have a body*. Then a skeptical hypothesis is raised: *her twin did it; I'm a disembodied spirit being deceived by an evil demon*. The hypothesis seems to undermine the knowledge claim. Part of the motivation for the relevant alternatives framework is that it provides an elegant, uniform diagnosis of these phenomena. When first made, the initial claims do count as knowledge (if true). This is because all the relevant alternatives are ruled out by the subject's evidence. But when attention is drawn to the skeptical hypothesis, it becomes salient and thereby relevant. Such hypotheses are, by design, alternatives to putative knowledge claims. So when they become relevant, they must be ruled out in order for the knowledge claim to stand in this new, stricter context. But they cannot be ruled out on the basis of our evidence, and so the claims no longer amount to knowledge: for all I can tell, the twin *could* have done it; maybe I *am* a disembodied spirit being deceived by an evil demon; and so on.

Now that I've set out the relevant alternatives framework, I can use it to shed light on the phenomenon of fake news.

2.2. Fake news and relevant alternatives

The core idea is this. When an agent, *S*, encounters a fake news story, it makes salient and thereby relevant a claim, *q*.⁵ Such a claim will be an alternative to all sorts of things *S* might know, $p_1 \dots p_n$. Thus the fake news story shifts *S* into a context where this is another relevant alternative to $p_1 \dots p_n$ that she must be in a position to rule out, in order to know those claims. But agents will not always, perhaps not even typically, have the evidence to rule out *q*. This threatens the agent's knowledge: she is now in a context where she knows, or is in a position to know, less than before. An example will help make this vivid.⁶

In 2016, then quarterback of the San Francisco 49ers, Colin Kaepernick, knelt in protest during the pre-game renditions of the national anthem. Kaepernick's protest was targeted particularly at police violence against black people in the United States. His actions sparked public uproar, and

⁵The term 'encounters' is supposed to cover various possibilities: reading, overhearing, being testified to, and so on.

⁶Unless otherwise referenced, details are taken from <https://www.snopes.com/fact-check/tim-tebow-kneel-anthem>.

he no longer plays in the NFL. This has been a politically important news story, because it serves as a proxy for various issues, including racial justice, employment fairness, and wealth inequality. One of the central questions is whether Kaepernick was treated fairly by the NFL. He alleges not, and has brought a collusion suit against the league.⁷ An important claim in assessing Kaepernick's actions, and the reactions to them, is the following.

KAEPERNICK. In 2016, Colin Kaepernick was the first quarterback to protest a political issue by kneeling during the anthem.

Most of both Kaepernick's supporters and detractors agree that KAEPERNICK is true. In the current context, this strong testimonial evidence puts you in a position to know that claim. So far, so good.

But, in 2018, a screenshot of another kneeling former NFL quarterback, Tim Tebow, did the rounds on various social media platforms. It was accompanied by the following text.

Let's not forget about Tim Tebow (who I respect very much), the NFL quarterback that kneeled in protest of abortion during the National Anthem in 2012 and was praised by fans for being a 'model American,' then was given 3 new NFL contracts in edition [*sic*] to a Major League Baseball contract. Unlike Colin Kaepernick, who CLEARLY [*sic*] stated his protest has nothing to do with the flag or military, Tim Tebow specifically said that he couldn't stand for a flag that allowed for abortion. If you can't see the hypocrisy in this then chances are you apart [*sic*] of the problem America has faced for centuries.

The story makes salient and thereby relevant the following claim.

TEBOW. In 2012, Tim Tebow protested a political issue by kneeling during the anthem.

TEBOW is an alternative to KAEPERNICK: they cannot both be true. Since it is now a relevant alternative, you must now be in a position to rule out TEBOW in order to know KAEPERNICK.

There are two reasons this is difficult. The first is that there is a gap between lacking sufficient evidence for believing TEBOW and being able to rule it out. I lack sufficient evidence for believing that the number of stars in the universe is even, but I cannot rule it out. So even if you don't believe the story, that doesn't mean you're thereby in a position to rule TEBOW out.

⁷See <https://www.nytimes.com/2018/08/30/sports/colin-kaepernick-collusion-case-nfl.html>.

The second reason that it's difficult to know KAEPERNICK in light of TEBOW's salience is that salience is contagious. A claim is rarely made salient in isolation. Rather, q 's salience sparks the imagination, making salient other claims. So even if you are in a position to rule out TEBOW, its salience carries in its wake the salience of other alternatives to KAEPERNICK, such as:

OTHER. Before Kaepernick protested a political issue by kneeling during the anthem, some other player did so.

OTHER is a weaker claim than TEBOW, and is thereby harder to rule out. So even if your evidence rules out the latter, the former's salience may nonetheless threaten your knowledge.⁸

These points generalise. Except in the most skeptical of contexts, an agent knows all sorts of things. Suppose in particular that she knows that p . By (RA), she is in a position to rule out all the relevant alternatives to p . But now she encounters a fake news story that threatens her knowledge of p in two ways. First, the story makes salient and thereby relevant a claim, q_1 , that is an alternative to p . But the agent is not in a position to rule out q_1 . By (RA), she does not know that p relative to this new set of alternatives. Crucially, the agent's knowledge is threatened even if she doesn't believe q_1 : it's enough that she is not in a position to rule it out. Call this a *direct threat* to the agent's knowledge.

Second, the fake news story can constitute instead (or additionally) an *indirect threat* to her knowledge. Such a story makes salient and thereby relevant a claim, q_2 , that is an alternative to p . This time the agent is in a position to rule out q_2 , so her knowing that p is not directly threatened. But salience is contagious. In making salient q_2 , the story makes salient and thereby relevant various other claims, $q_3 \dots q_n$. If at least one of these claims is both an alternative to p and such that the agent is not in a position to rule it out, then she is now in a context where she does not know that p .

Before moving on, I need to consider an objection. If our knowledge is so easily threatened according to the relevant alternatives framework, then so much the worse for that framework. Consider this case:

OBAMA. You're at the bus station and overhear someone claiming that Barack Obama was born in Argentina. He supports this claim with a story about the

⁸I don't know the exact provenance of the Tebow story, but it is likely that it is a piece of fake news in the sense defined in §1. Does that piece of information now put you in a position to rule out TEBOW? That is a tricky question, to which I'll return in §4.

ease of faking Hawaiian birth certificates. You don't know much about birth certificates, so it seems you're not in a position to rule out that Obama was born in Argentina.

According to (RA), the objection continues, after overhearing a stranger at the bus station, you no longer know that Obama was born in the USA. This is because now the alternative that he was born in Argentina is relevant and you cannot rule it out. It seems wrong that such an encounter threatens your knowledge. Moreover, these kinds of situations arise all the time. But it's absurd to think that our knowledge is constantly under threat in this way; this is a *reductio* of (RA).

My response is twofold. First, the objection's intuitive force derives from the particular knowledge claim it invokes and an underspecification of the case. The case is *prima facie* problematic because it seems to put our knowledge at the mercy of epistemically malevolent agents, leaving us at a loss to know many important claims. However, there has been much public discussion around the potentially known claim involved in the case, that Obama was born in the USA. Given this, most agents *will* be in a position to rule out the claim that he was born in Argentina. They will have lots of evidence that rules out that alternative: it has been debunked by numerous, reputable investigative sources; there is no governmental coverup about Obama's birthplace; presidential candidates' credentials are thoroughly vetted; and so on. So even though an alternative is now relevant, most agents will still be in a position to rule it out and so will know that Obama was born in the USA.⁹

Suppose we flesh out the case in the other direction. Suppose you know before overhearing the stranger that Obama was born in the USA, but you have very little evidence to support this claim. You don't even remember that Obama was President, or you don't remember that presidents must be born in the USA.¹⁰ Then I agree that, after hearing the stranger, you no longer know that Obama was born in the USA. But this is just as it should be. After all, in this version of the case you don't have any independent evidence to rule out the stranger's claim. It would be very implausible to let an agent rule out *q*, an alternative to *p*, on the basis of knowing *p* itself (relative to some other set of alternatives that does not include *q*).

⁹Does this undermine the problem that fake news is supposed to raise? No. Lots of fake news stories, like the one about Tebow, make salient alternatives to potential knowledge claims that many agents are not in positions to rule out.

¹⁰One might think this is too scant a basis on which to know that Obama was born in the USA. If so, the present objection doesn't get off the ground, since the problem is supposed to be that knowledge is too easily lost.

That would mean knowledge could never be lost once gained, because one could always use p itself to rule out any competing alternative.

Second, suppose instead that the objection is that, in general, knowledge shouldn't be so easily lost. The sense that this is problematic is largely defanged by specifying OBAMA in more detail; similar strategies will work for similar cases. An alternative threatens knowledge only if it is relevant *and* cannot be ruled out. Agents will lose knowledge when and only when they cannot rule out the alternative that has become relevant. Any lingering dissatisfaction may just come down to a clash of intuitions. This version of the objection rejects the idea that knowledge that p – whatever p might be – can be lost and gained as easily as an alternative can become relevant or not. But many see this commitment as a strength, especially when diagnosing and dealing with skeptical challenges. I also think it usefully illuminates the epistemic problems posed by fake news.

One of those problems is that fake news directly or indirectly threatens agents' knowledge claims, by making salient and thereby relevant alternatives that they may not be in a position to rule out. Call this *the problem of relevant alternatives*.

There's another problem that I want to highlight here. The problem of relevant alternatives is a micro-epistemological problem, in the sense that it occurs in an individual's encounters with pieces of fake news. The second problem is a macro-epistemological problem, arising in larger contexts of whole epistemic economies. Let me explain.

3. The degradation of our epistemic environment

3.1. Epistemic environments

Consider an epistemic community: a collection of agents who go about the world knowing, believing, testifying, and so on. Whether a member of the community is in a position to achieve some positive epistemic status – knowledge, understanding, justified belief, and so on – depends on a variety of things: the totality of resources and circumstances relevant to assessing epistemically interesting statuses. Call the totality of these things the community's *epistemic environment*.¹¹ Fully characterising the notion of an epistemic environment is beyond the present scope. Instead, let me ostend the concept by highlighting some of its prominent features.

¹¹The phrase can be found in Levy (2018), though he leaves its characterisation completely intuitive.

A community's epistemic environment includes the community's experts on matters various. Unless I'm a member of a community with, say, expertise in particle physics, my epistemic environment is not conducive to my having knowledge about fundamental particles.¹² Reflection on this kind of case makes clear that a community's epistemic environment will also include its technological resources: knowledge claims about bacteria are hard to come by without microscopes. Conceptual and other hermeneutical resources also contribute to an epistemic environment; we can usefully construe hermeneutical injustice as an agent's being in an epistemic environment insufficient for understanding their own experience (Fricker 2007). An epistemic environment also involves more obviously environmental factors. Knowledge and its kin will be hard to gain in worlds where nature is far from uniform. Finally, though by no means exhaustively, an epistemic environment will include various things a member of the community is in a position to know, or at least rationally believe, about the environment itself. If the environment is such that an agent knows that lying is frequent, it will be much harder for her to gain knowledge by testimony.

An epistemic environment is relative to an epistemic community, and it changes as the circumstances and resources of that community change. But it is more stable than the contexts that relevant alternative theorists take to be important for knowledge. Contextual features that determine which alternatives are relevant are very fine-grained and in constant flux, shifting as easily as alternatives flit in and out of salience. Epistemic environments do not shift in such radical ways from moment to moment.

Epistemic environments can be better or worse, insofar as they are more or less conducive to members of the epistemic community achieving positive epistemic statuses.¹³ The actions and speech acts that we perform, both individually and collectively, as well as the character traits that we manifest, can thus ameliorate or degrade our epistemic environment. Making a reliable resource widely available might constitute a significant improvement to a community's epistemic environment, as might cultivating the virtues of honesty, openness, and humility; spreading misleading evidence of falsehoods might make for an impairment.

¹²For ease of exposition, I'll use knowledge to stand for all the various epistemically interesting statuses that we might be concerned with.

¹³A complication that I mention to set aside: an epistemic environment may be assessed along different dimensions, according to different epistemic statuses. For instance, consider an environment in which you're deceived by a strange demon. They make the world appear such as to enable you to garner all sorts of rational beliefs, though knowledge is extremely difficult to come by.

3.2. Fake news and the degradation of the epistemic environment

With the notion of an epistemic environment in hand, we can see that fake news poses a problem in addition to the problem of relevant alternatives outlined in §2.2. The problem is that it degrades our epistemic environment: it makes it harder for us to achieve positive epistemic statuses. I'll focus on the achievement of knowledge, and on three related ways in which fake news degrades our epistemic environment in this respect.

First, is the macro version of the micro problem of relevant alternatives. The more fake news stories that are in circulation, the more alternatives are made salient and thereby relevant – alternatives that agents must be in a position to rule out.

Second, is the deterioration of confidence in our epistemic institutions. It's plausible that repeated exposure to fake news claims can undermine confidence in epistemic institutions, like testimony.¹⁴ This undermining constitutes a degradation of the epistemic environment insofar as it makes transmission of knowledge by testimony – and thereby the achievement of knowledge in general – more difficult. If agents are not confident in their community's epistemic institutions, it will be harder for them to gain positive epistemic statuses.

Third, is altering of epistemic habits. This is a natural response to the first two points. The confrontation with fake news might train agents to consider wider ranges of alternatives when evaluating claims, even when they haven't been directly presented with these alternatives.¹⁵ And agents might internalise their lack of confidence in epistemic institutions, being generally less willing to believe testimony given their undermined confidence in it. This point differs from the second in that no change of habit need follow the undermining of confidence in an epistemic institution. We may explicitly tell ourselves that social media testimony is not to be trusted, and yet find ourselves responding to it just as before.

It's important to note that the second and third points – the undermining of confidence in our epistemic institutions and the altering of epistemic habits – degrade the epistemic environment even if they are rational or good responses to the proliferation of fake news. Better

¹⁴It would be good to have empirical support for this claim, but I'm not aware of anything available other than anecdotal evidence.

¹⁵This might help to explain why some people are prone to skepticism about official versions of events. Their habit of generating alternatives to official explanations makes it harder for them to know those explanations. This point is neutral on the issue of whether such habits are, following Cassam (2016), to be seen in general as intellectually vicious or, following Coady (2003, 2012) and Pigden (2017), as epistemically called for in many environments, including our own.

epistemic environments are more conducive to positive epistemic statuses, like knowledge. It may be no fault of an agent's that she is in a poor epistemic environment. Moreover, given that she is, it may be permissible or even obligatory that she withhold judgment in various ways. In the extreme case, the person who gets overwhelming evidence that they're being deceived by an evil demon might have to suspend belief about almost everything. That might be the best they can make of a bad situation; nonetheless their environment is still highly impoverished with respect to the acquisition of knowledge.

This relates to Rini's (2017) contention that partisan epistemology can be epistemically virtuous in epistemically non-ideal settings. Very briefly, her argument is this. Partisanship is a proxy for a variety of normative judgments and commitments. So if you learn that someone shares a political affiliation with you, that indicates that they share views with you about a variety of normative matters, which licenses giving them a credibility boost on what she calls 'politically relevant claims'. Rini goes on to develop these points into an argument to the effect that solutions to the problems raised by fake news should be institutionally rather than individually focused. She gives the example of social media platforms tracking their users' credibility: habitual sharers of fake news would still be able to post, but there'd be a red flag by their name, or something similar.

I agree that institutional solutions to fake news will be important. And I agree that in epistemically non-ideal settings, individual epistemic practices may be excused that would not be condoned in ideal settings.¹⁶ But I think Rini underemphasises the importance of holding individuals responsible for their contribution to the degradation of the epistemic environment. Doing so is an important part of trying to solve the problem; I turn to this now.¹⁷

¹⁶Note that this is weaker than the conclusion Rini argues for, namely that partisan epistemology can be virtuous, or reasonable, or justified with respect to fake news. Moreover, while I think many agents may be excused for propagating fake news, I don't require partisan epistemology to be the exculpating factor. More on this in §4.1.

¹⁷A word about how my discussion of fake news meshes with Levy's (2017). On his view, the main problem posed by fake news stems from our individual psychological constitutions. He brings to bear recent work in cognitive science on how agents' beliefs are self-ascribed, and how they persist in the face of countervailing evidence, and argues that fake news stories are likely to have similar features. In particular, even when representations are stored explicitly not as beliefs – even as false or fictional – they might later be self-ascribed as beliefs, or might nonetheless have effects on agents' behavior. This is because, Levy argues, representations are not stably and systemically tagged with attitude flags – belief, imagination, and so on. Rather, the attitude ascriptions that representations receive are malleable. Thus even though an agent represents p initially as part of a fictional narrative, she might later treat p as the content of a belief, either in self-ascription or in behavior. This is all compatible

4. How to proceed

I've raised two problems that fake news poses. I'll end by sketching some potential responses and defending a solution that recommends the social condemnation of those who spread fake news.

4.1. *The problem of the degradation of the epistemic environment*

The problem is that fake news degrades our epistemic environment. By undermining our confidence in epistemic institutions, and by altering our epistemic habits, it makes our environment less conducive to achieving positive epistemic statuses such as knowledge.

While I agree with Rini (2017) that institutional changes are important, I think bottom-up, individual-level change is also needed.¹⁸ Top-down change without bottom-up support is rarely very effective.¹⁹ 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.

What might this bottom-up reproach look like? At a minimum, we should criticise creators and distributors of fake news. Social sanction is an effective method of norm enforcement and internalisation. For instance, we should be less willing to let purveyors of fake news get away with the 'a retweet is not an endorsement' defence that Rini (2017) highlights is common. As I've argued, fake news degrades the epistemic environment even when it is shared without malice or mischief. To effect the large shift in testimonial practices that tackling fake news requires, we'll need to treat infractions with more seriousness.

One might object at this point that we can't reproach individuals for their engagement with fake news. After all, Rini (2017) argues that believing fake news can be individually reasonable in certain epistemically non-ideal settings. In particular, given that partisanship is a proxy for normative discernment, it can be reasonable to give extra credence to claims made by partisans. If an individual comes to believe a piece of fake

with the account I've given. I read Levy as interested primarily in expounding the cognitive mechanisms that fake news might problematically exploit, rather than its more general normative assessment.

¹⁸To be clear, Rini does not claim that it is *not* needed. We may not disagree much at all about how to tackle fake news, but to the extent that we do, it is primarily over emphasis. See especially Rini (2017, 55).

¹⁹This is a point that some feminist theorists, such as Bartky (1990), have argued for in the context of dominance and norms of oppression. However, this is a point of active contention in the feminist literature. For recent top-down approaches, see Young (2011) and Haslanger (2015).

news from a partisan source, then it seems they may well be epistemically blameless.

I think Rini's argument overreaches. Even if she is correct that it can sometimes be permissible to give credibility boosts to epistemic partisans, it is not clear that it is reasonable or justified to do this in contexts where fake news is common, like social media.²⁰ Spreading fake news and thereby contributing to the degradation of the epistemic environment is an epistemic wrong.

There's a nearby challenge, however. It's not obvious how to deal with fake news, and many people seem to be unwittingly guilty of propagating it. So even if individuals are not justified in their spreading of fake news, they may be *excused* for doing so. Hence, while I'm not entirely convinced that propagators of fake news are epistemically justified, there's a strong case for thinking that they're excused. And that raises a question. How can we reproach blameless propagators of fake news?

How to answer this question depends on how one thinks about the relationship between responsibility, reproach, and wrongdoing. I will draw on Calhoun's (1989) understanding of responsibility and reproach to respond to this challenge. Before doing so however, I'll sketch two other responses for readers whose commitments differ from my own.

The first response is to accept the claim that believing fake news is excusable, but deny the corresponding claim about *reproducing* it. This concedes that we should not reproach individuals for being taken in by fake news, in many cases, but we nonetheless can reproach them for spreading fake news. There's a tension in this response, however. Beliefs are supposed to be relied upon in future deliberation and action. There's something unsatisfyingly Janus-faced about excusing the acquisition or sustenance of a belief but condemning acting on it, for instance by sharing the claim on social media. It seems odd to allow the acquisition of a belief to be epistemically blameless, but reproducing its content epistemically blameworthy. Nonetheless, perhaps this oddness is something we have to live with, given the non-ideal situation that we find ourselves in.

The second response builds on the distinction between culpability, blameworthiness and reproach on the one hand, and responsibility and liability on the other. A variant of a case due to McMahan (2005, 393) brings out the difference.

²⁰For arguments against Rini's defence of epistemic partisanship, see Worsnip (2019).

DRIVER. Alma keeps her car well maintained and is an impeccable driver. One day, due to an unfortunate and unlikely mechanical failure, the brakes and steering on Alma's car fail and she hits and kills a pedestrian.

McMahan's judgment is that Alma is not culpable for the death of the pedestrian. After all, she did everything in her power to keep her car in good condition, and was as good a driver as any. Nonetheless, she is responsible for the pedestrian's death. She knowingly got behind the wheel of a dangerous machine, aware that there were risks of wrongdoing, even if she tried her best to avoid them.

Although DRIVER was deployed in the context of self-defense, we might apply a similar lesson to the case of believing and spreading fake news. Even if many agents have excellent excuses and so are not culpable for the degradation of the epistemic environment, they are nonetheless responsible for their part in that degradation. They knowingly engage in risky epistemic practices: accepting testimony. 'Risky' here needn't involve a high level of risk. Alma's risk of causing a fatal accident, given her car's service record and her driving skill, may be extremely low. Similarly, an epistemic agent may be fairly good at discriminating true claims from false, at not being taken in by fake news stories, and so on. Nonetheless, she takes an epistemic risk by engaging in the discourse and so can be held responsible when things go wrong.

Both of these responses have costs. The first creates an implausible asymmetry in the evaluation of belief acquisition and dissemination. The second cannot underwrite the social criticism that I think we need in order to shift our practices. Even if we hold Alma responsible for killing the pedestrian, she is not to be reproached for anything she does. We can completely condone her behavior and recommend others to act just as she did. But that is not going to help with fake news. The goal with fake news, unlike DRIVER, is to get agents to change their behavior, or perhaps to internalise new norms. Holding them responsible without reproaching, criticising, or otherwise blaming them is not going to effect the changes we need to combat fake news in our epistemic communities.

Because of this, I prefer a different response to the challenge of reproaching agents who might nonetheless be epistemically blameless. In the context of moral responsibility and culpability for the oppression of women, Calhoun (1989, 393–398) distinguishes between *normal* and *abnormal moral contexts*. In normal moral contexts, there is widespread knowledge of, and compliance with, correct moral standards. In abnormal

moral contexts, ignorance and noncompliance with correct standards is commonplace, usually because the community's standards are incorrect. For that reason, Calhoun (1989, 400–405) argues, it is hard to excuse nonculpable violators in abnormal contexts without thereby condoning their behavior. The excusing conditions – ignorance of the correct norms; that everyone else is doing it – also serve to justify the actions. Calhoun (1989, 402) imagines the response of a male scientist being brought to task for being problematically unreflective in his academic practices; the critic acknowledges that he is excused because everyone else does it: '[w]hat will his response be? "That's not an excuse, that's my reason!"'. Calhoun's conclusion is that whereas it is not appropriate to reproach those who are blameless in normal moral contexts, it may well be appropriate to do so in abnormal moral contexts. This is because reproaching nonculpable wrongdoers in such situations is required to avoid condoning the wrongdoing.

Calhoun's framework can be usefully recruited to solve the present problem. The challenge is to see how one could reasonably reproach nonculpable purveyors of fake news. On the one hand, it is important to deploy social sanction and criticism to combat the degradation of the epistemic environment. On the other hand, it seems inappropriate to blame the blameless. But we can draw a parallel distinction to Calhoun's, a distinction between normal and abnormal *epistemic* contexts. In normal epistemic contexts, there is widespread knowledge of and compliance with the correct epistemic standards. In abnormal epistemic contexts, ignorance and noncompliance is commonplace.

We are in an abnormal epistemic context with respect to fake news. Our current testimonial practices do not seem up to task, and there isn't widespread consensus on how to improve them. Calhoun's framework predicts that it will be hard in this context to point out nonculpable wrongdoing without condoning it. And this is just right. Suppose you tell someone who unwittingly propagates fake news that they committed an epistemic wrong, but are excused because they couldn't have known any better, or because almost everyone else is doing the same thing. This might well provoke the same reaction as Calhoun's scientist: that's a justification, not an excuse!

However, we should censure violations of good epistemic practice that contribute to the degradation of the epistemic environment. If the abnormality of our epistemic context makes this hard to do blamelessly, then we should deploy blame. We should reproach even the nonculpable propagators of fake news. If we do not, we deprive ourselves of a crucial

tool in the regulation, enforcement, and change of standards: social condemnation. And in so depriving ourselves, we will make it much harder to shift our epistemic practices in the ways that we need.

4.2. The problem of relevant alternatives

Whereas the problem of the degradation of the epistemic environment is a macro problem, concerning entire epistemic communities, the problem of relevant alternatives is a micro problem, found at the level of individual epistemic transactions. The problem is how to combat the threat that fake news stories pose to knowledge claims. As I explained in §2.2, the problem arises because such stories make salient and thereby relevant alternatives to knowledge claims that agents may not be in a position to rule out.

There are two complementary ways to go. The first is to retreat to a context where the alternative is no longer relevant, and so does not need to be ruled out. I think we currently lack the epistemological resources to say much of interest about how this might look. The relevant alternatives literature has arisen largely to diagnose classical skeptical problems. Lots of focus has been given to how alternatives become relevant, but little has been given to how they lose relevance. After all, having guided oneself through Descartes's *Meditations*, it is not as though the evil demon hypothesis is forevermore a relevant alternative to most of one's knowledge claims.

It's not clear how an alternative loses relevance, but one suggestion for further research is the following. Even though it is easy to make an alternative relevant, not all alternatives stay relevant in the same way. As mentioned, absent explicit consideration, radical skeptical hypotheses fade from relevance quickly. Let's say that they are not *robustly* relevant alternatives. Keeping track of a wide variety of relevant alternatives is both cognitively demanding and conducive to skepticism. So perhaps individuals are furnished with a default range of relevant alternatives – a range of plausible hypotheses that might compete with potential knowledge claims. Alternatives in this range are robustly relevant, needing to be ruled out across many contexts in order for knowledge. Alternatives outside the range are not robustly relevant: they can be made salient, but will quickly fade from relevance if not frequently reconsidered.

An individual's epistemic environment will no doubt contribute to which alternatives are and are not robustly relevant. So one way to make progress might be to shape our epistemic environments so as to

preclude likely fake news claims from robust relevance. This might happen by source – claims from such-and-such a testifier are not to be part of the default stock of alternatives – or by content – claims about so-and-so topic are not to be part of the default stock either.

The second way to tackle the problem of relevant alternatives is to give agents evidence to rule out the fake news alternatives. Then, given that we may have little control of whether the fake news alternative is relevant or not, agents will still be in a position to rule it out and so it won't threaten their knowledge even when it is relevant. This will partly involve providing good evidence when raising fake news alternatives to debunk them; otherwise there is a risk of the debunking strategy backfiring. By drawing attention to q , even as fake news, you might thereby make it a relevant alternative to your interlocutors' knowledge claims which they are not in a position to rule out. It will also partly involve improving general epistemic literacy, by making agents aware of good debunking resources such as Snopes.²¹

Unfortunately, this strategy may be of limited efficacy. As I argued in §2, salience (and thereby relevance) is contagious. So even if you provide interlocutors with evidence against q , you can never be sure of outfitting their evidential armory with evidence to rule out other alternatives to their knowledge, the salience of which is parasitic on q . For instance, a privately flown drone supposedly shutdown Gatwick airport in the United Kingdom in December 2018.²² Suppose you try and debunk the claim that the object seen was really a bird. You may succeed in furnishing me with evidence that the object was not a bird (its flight pattern was inconsistent with those of any local bird species, say), but now the alternative that the object was something *else* – a plastic bag, police surveillance equipment, and so on – has been made salient and thereby relevant.

The upshot is pessimistic; our knowledge is easily undermined by fake news narratives. So perhaps prevention is the best strategy of all. Agents should be counselled not to share stories unless they've ruled out that the story they want to share is fake news.

5. Conclusion

I raised two new problems posed by the rise of fake news. The first occurs at the microlevel. The problem of relevant alternatives is that fake news

²¹<http://www.snopes.com>

²²<https://www.nytimes.com/2018/12/27/world/europe/gatwick-airport-drone.html>

stories make salient and thereby relevant alternatives that threaten our knowledge claims. The second problem arises at the macrolevel. The problem of the degradation of the epistemic environment is that the proliferation of fake news makes our epistemic environment less conducive to achieving positive epistemic statuses, such as knowledge and understanding. It does this primarily by eroding confidence in epistemic institutions and by altering epistemic habits.

What to do about all of this is a hard question, and there is not much reason to be optimistic. But it's important to stress personal and interpersonal accountability, in addition to institutional solutions that have already been offered. I doubt we'll make much progress on the problems that fake news poses unless we start taking it seriously in our epistemic communities, by reproaching those that spread it.

This is a hard pill to swallow because believing and spreading fake news may be nonculpable in our epistemic context. But the abnormality of the context makes it important to reproach one another for believing and spreading fake news, even when we do so faultlessly.

Acknowledgments

This paper grew out of a presentation I gave at UNC's 'Philosophy in 15 Minutes' event. I'm grateful to Mariska Leunissen for the invitation to present and so the opportunity to think about these issues. Thanks also to Joanna Lawson, Silvan Wittwer, and Alex Worsnip for helpful input at various stages of this project. Finally, thank you to two anonymous reviewers, whose generous comments improved the paper greatly.

Disclosure statement

No potential conflict of interest was reported by the author(s).

ORCID

Christopher Blake-Turner  <http://orcid.org/0000-0002-1509-0649>

References

- Aikin, Scott F., and Robert B. Talisse. 2018. "On 'Fake News'." 2018. <https://www.3quarksdaily.com/3quarksdaily/2018/05/on-fake-news.html>.
- Bartky, Sandra Lee. 1990. *Femininity and Domination: Studies in the Phenomenology of Oppression*. New York: Routledge.

- Calhoun, Cheshire. 1989. "Responsibility and Reproach." *Ethics* 99 (2): 389–406. doi:10.1086/293071.
- Cassam, Quassim. 2016. "Vice Epistemology." *The Monist* 99 (2): 159–180. doi:10.1093/monist/onv034.
- Coady, David. 2003. "Conspiracy Theories and Official Stories." *International Journal of Applied Philosophy* 17 (2): 197–209. doi:10.5840/ijap200317210.
- Coady, David. 2012. *What to Believe Now: Applying Epistemology to Contemporary Issues*. Chichester: Wiley-Blackwell.
- Cohen, Stewart. 1988. "How to Be a Fallibilist." *Philosophical Perspectives* 2: 91–123. doi:10.2307/2214070.
- DeRose, Keith. 1992. "Contextualism and Knowledge Attributions." *Philosophy and Phenomenological Research* 52 (4): 913–929. doi:10.2307/2107917.
- DeRose, Keith. 1995. "Solving the Skeptical Problem." *Philosophical Review* 104 (1): 1–52. doi:10.2307/2186011.
- Dretske, Fred I. 1970. "Epistemic Operators." *Journal of Philosophy* 67 (24): 1007–1023. doi:10.2307/2024710.
- Frankfurt, Harry. 2005. *On Bullshit*. Princeton: Princeton University Press.
- Fricker, Miranda. 2007. *Epistemic Injustice: Power and the Ethics of Knowing*. New York: Oxford University Press.
- Gelfert, Axel. 2018. "Fake News: A Definition." *Informal Logic* 38 (1): 84–117. doi:10.22329/il.v38i1.5068.
- Haslanger, Sally. 2015. "Distinguished Lecture: Social Structure, Narrative and Explanation." *Canadian Journal of Philosophy* 45 (1): 1–15. doi:10.1080/00455091.2015.1019176.
- Hawthorne, John. 2003. *Knowledge and Lotteries*. Oxford: Oxford University Press. doi:10.1093/0199269556.001.0001.
- Levy, Neil. 2017. "The Bad News about Fake News." *Social Epistemology Review and Reply Collective* 6 (8): 20–36.
- Levy, Neil. 2018. "Taking Responsibility for Health in an Epistemically Polluted Environment." *Theoretical Medicine and Bioethics* 39 (2): 123–141. doi:10.1007/s11017-018-9444-1.
- Lewis, David. 1996. "Elusive Knowledge." *Australasian Journal of Philosophy* 74 (4): 549–567.
- McMahan, Jeff. 2005. "The Basis of Moral Liability to Defensive Killing." *Philosophical Issues* 15 (1): 386–405. doi:10.1111/j.1533-6077.2005.00073.x.
- Mukerji, Nikil. 2018. "What Is Fake News?" *Ergo, an Open Access Journal of Philosophy* 5 (35). doi:10.3998/ergo.12405314.0005.035.
- Nagel, Jennifer. 2008. "Knowledge Ascriptions and the Psychological Consequences of Changing Stakes." *Australasian Journal of Philosophy* 86 (2): 279–294. doi:10.1080/00048400801886397.
- Pepp, Jessica, Eliot Michaelson, and Rachel Katharine Sterken. Forthcoming. "What's New About Fake News?" *Journal of Ethics and Social Philosophy*.
- Pigden, Charles R. 2017. "Are Conspiracy Theorists Epistemically Vicious?" In *A Companion to Applied Philosophy*, edited by Kasper Lippert-Rasmussen, Kimberley Brownlee, and David Coady, 120–132. Chichester: Wiley-Blackwell.

- Rini, Regina. 2017. "Fake News and Partisan Epistemology." *Kennedy Institute of Ethics Journal* 27 (2): E-43–E-64. doi:10.1353/ken.2017.0025.
- Rysiew, Patrick. 2001. "The Context-Sensitivity of Knowledge Attributions." *Nous* 35 (4): 477–514. doi:10.1111/0029-4624.00349.
- Rysiew, Patrick. 2007. "Speaking of Knowing." *Nous* 41 (4): 627–662. doi:10.1111/j.1468-0068.2007.00665.x.
- Stine, Gail C. 1976. "Skepticism, Relevant Alternatives, and Deductive Closure." *Philosophical Studies* 29 (4): 249–261. doi:10.1007/bf00411885.
- Stoutenburg, Gregory. 2017. "Strict Moderate Invariantism and Knowledge-Denials." *Philosophical Studies* 174 (8): 2029–2044. doi:10.1007/s11098-016-0786-0.
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.
- 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*, edited by Carl Fox, and Joe Saunders, 240–264. London: Routledge.
- Young, Iris Marion. 2011. *Responsibility for Justice*. Oxford: Oxford University Press.