

UC Irvine

UC Irvine Previously Published Works

Title

Good News, Bad News, Fake News

Permalink

<https://escholarship.org/uc/item/3nw982tc>

Author

Pritchard, Duncan

Publication Date

2021-06-10

DOI

10.1093/oso/9780198863977.003.0003

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NoDerivatives License, available at <https://creativecommons.org/licenses/by-nd/4.0/>

Peer reviewed

For *Epistemology of Fake News*,
(eds.) S. Bernecker, A. Flowerre & T. Grundman, (Oxford: Oxford University Press).

GOOD NEWS, BAD NEWS, FAKE NEWS

DUNCAN PRITCHARD

University of California, Irvine & University of Edinburgh

ABSTRACT. An account is offered of the nature of fake news, and it is explained how this account differs from the main proposals in the contemporary philosophical literature in this regard. One key feature of the account is the idea that fake news is not a genuine form of news. In particular, fake news is to be distinguished from genuine news that is epistemically problematic. It is argued that this point is important because it entails that what is required to differentiate news with a sound epistemic pedigree from news that has a poor epistemic pedigree is distinct from what is required to differentiate genuine news from fake news. This has implications for how we should manage the challenge posed by fake news, at both the individual and the structural levels.

*“All we have to do now/
is take these lies/
and make them true/
somehow.”*

George Michael, *Freedom*

0. INTRODUCTORY REMARKS

The contemporary terminology of ‘fake news’ is largely a term of art. Pundits, politicians, journalists and such like use this terminology in an undisciplined way, as one would expect. In particular, it is often used as an insult to throw at one’s enemies. As such, it can sometimes be rather misleading, in that something might be labelled ‘fake news’ when its only crime is to provide information that some people find contrary to their interests.¹ Naturally, we wouldn’t want an account of fake news that took this usage at face value (which is not to say, of course,

that it is irrelevant to this account why this terminology is used in this way at all), since it would make the concept so broad as to be empty. For one thing, it would entail that the concept could be legitimately applied in contradictory ways, such that the very same phenomenon can be rightly classified as both fake news and not fake news. Accordingly, if we want to make some philosophical headway understanding this phenomenon, it's important that we first do a little conceptual ground-clearing.

The kind of account of fake news that I want to defend is one on which fake news is to be contrasted with genuine news (including genuine news that has a poor epistemic pedigree), in terms of how it involves deliberately conveying misleading information with an intent to mislead. Fake news is presented as news, but it is not a kind of news at all (not even an epistemically deficient form of news), any more than an excellent forgery is thereby the real thing. Note that on my view fake news needn't involve the presentation of falsehoods (though it typically will), given that even the literal truth can be misleading (indeed, this is sometimes the most effective way to mislead). In addition, while fake news involves an intent to mislead on my proposal, it doesn't follow that if it doesn't mislead then it isn't fake news; ineffective fake news is still fake news.

Aside from the intrinsic interest in having the right account of fake news in hand, there are also practical implications to this account. In particular, as we will see, this account of fake news has ramifications for how we might go about dealing with the challenges posed by fake news, at both an individual and structural level.

1. UNDERSTANDING FAKE NEWS

In order to understand this proposal, it will be helpful to contrast it with the most developed account of fake news in the philosophical literature, due to Axel Gelfert (2018), not least because this account seems very similar to the view I just outlined.² After a comprehensive discussion of different kinds of fake news, and different treatments of this notion, both philosophical and non-philosophical, he ends up with the following definition of fake news: “the deliberate presentation of (typically) false or misleading claims *as news*, where the claims are misleading *by design*.” (Gelfert 2018, 108, *italics in original*)³ On the face of it, Gelfert seems to be endorsing a very similar account to the one that I offer. As we will see, however, there

are in fact important divergences that we need to get a handle on.

Let's begin with what Gelfert gets right. Gelfert's account of fake news has the virtue of focusing on the intention and systemic goals of the fake news rather than on the veracity of the output. In particular, that the presentation of false or misleading claims as news must be deliberate, with the goal of misleading (i.e., this is 'by design'), is clearly right. Unintentional errors in a news report don't make it fake news, even though this would clearly be a case of news that involves falsehoods (and which might thereby mislead). Moreover, a good satirical publication like *The Onion* or the UK's *Private Eye* is not fake news, even if sometimes gullible people are taken in by their fake headlines, since it is not designed to mislead. Gelfert is also right that it's important that fake news is presented as news. As he notes (Gelfert 2018, 110), an advertising campaign could contain falsehoods and be designed to mislead, but that wouldn't thereby make it fake news. Putting the point in terms of a disjunction between false or misleading claims is also helpful, since sometimes the literal truth can be misleading, as when it leaves out important qualifiers or context. (Indeed, the most effective forms of fake news might well involve no literal falsehood at all).⁴ Note, though, that since false claims are also misleading, we can exclusively focus on the second disjunction, as I do in the summary of my proposal above, without loss of explanation (even though there might be a wider dialectical benefit to making it explicit that 'misleading' here also covers explicitly false claims).⁵

Nonetheless, despite its many merits, Gelfert's way of thinking about fake news is not quite right, and the reasons why are important, since they reveal deeper conceptual confusions. As we will see, the problem isn't quite the formulation above, but rather how Gelfert is unpacking this formulation. Moreover, understanding this point will in turn help us to get a handle on what the account of fake news I am proposing involves.

The first issue with Gelfert's proposal concerns the idea, which we have just endorsed, that the claims at issue in fake news are misleading by design. We have naturally unpacked this idea as meaning that it must be a goal of the fake news to mislead. As we noted, this feature of fake news would explain why a satirical news programme or publication wouldn't qualify as fake news. Note that it is consistent with our characterization of this clause that the agents involved have other motives in play, and even that these other motives are more central to their overall interests (since it is merely *a* goal of the fake news, and not *the* goal, or the overarching goal).⁶

This isn't the rendering of this clause that Gelfert has in mind, however. Instead, he explicitly argues that while fake news is misleading by design, it needn't be designed to mislead. That is, what is important to fake news on Gelfert's account is that it intentionally employs misleading information, but it needn't be the goal of the fake news that it has the effect of misleading its audience (even though, as we will explain in a moment, on Gelfert's view fake news must *actually* mislead its audience). The interesting example he gives to illustrate this point is of putatively fake news that is designed as 'clickbait'. There are websites whose goal (usually for financial reasons) is to encourage visitors, and Gelfert claims that often fake news is created for the purposes of clickbait. Presenting misleading claims as news can make effective clickbait, but since the goal of clickbait is just to generate a high volume of internet traffic, and not specifically to mislead, Gelfert argues that we should not build an intention to mislead into the definition of fake news.

While this is an intriguing proposal, I don't think that it is plausible on closer inspection. For consider what would be involved in developing a website that presented misleading claims as news purely with the goal of generating internet traffic, with no concern at all for whether the intended audience are actually taken in by these claims (i.e., are actually misled). Here is the rub: if the audience don't find the claims even remotely plausible, then why on earth would they be clicking on them under the guise of them being news? The point is that clickbait that wants to generate internet traffic by offering misleading claims *as news* must be in the business of also wanting to actually mislead its clientele, since that is precisely what is going to ensure that this strategy is successful. It is true that those who develop clickbait (*qua* fake news) will have other motives in play in addition to intentionally misleading people, but the point is that they must at least have this motive in play for their activities to be coherently pursued. Remember that our account of fake news only demands that it should be intentionally misleading; it doesn't require that this should be the only motive in play, or even that it should be the overarching motive in play.

Of course, we should note that there is a variety of clickbait that simply involves presenting outlandish claims, as there is a recognition that a certain kind of person is likely to click on the site as a result. But such sites do not offer any pretense of being news sources, and so are not even in the market for being considered fake news (including by the lights of Gelfert's account). There is thus no plausible sense to the idea that there can be fake news—i.e., claims that are presented as news that are misleading by design—which doesn't also

involve being designed to mislead. Either clickbait is presented as news, in which case it fits our rubric for fake news, or it isn't, in which case it is not a plausible candidate for fake news anyway, and so there is no surprise that it fails to satisfy the rubric.

Now one kind of case that one might think would work in Gelfert's favour in this regard is *bullshit*. As Harry Frankfurt (2005) has persuasively argued, what characterizes bullshit is a complete lack of concern for the truth. This sets bullshit apart from lying, as the liar usually does care about the truth, it's just that they want to hide it from the person they are lying to.⁷ In this sense the bullshitter will often be the source of misleading information, but not because her overarching goal is to deceive her audience, but rather because her lack of concern for the truth means that she is sanguine about bearing witness to such epistemically problematic reports. Imagine now that the bullshitter presents misleading information as news. Wouldn't this be a natural case of fake news? If so, given that the bullshitter doesn't have any overarching goal to deceive her audience, doesn't that speak in favour of Gelfert's account?

Once we examine this case more closely, however, it becomes clear that rather than supporting Gelfert's view, it in fact counts against it. The key thing to note is that the bullshitter's lack of concern for the truth doesn't entail that she is unconcerned with whether her audience believes what she says. On the contrary, the bullshitter wants to be believed. What distinguishes the bullshitter is rather that it doesn't matter to her whether what she asserts is something that she believes, and thus regards as true.⁸ Indeed, her unconcern for the truth will mean that she would rather assert something expedient to her interests than something she believes to be true. While she doesn't have any overarching concern to deceive others, that's only because if it is useful to her to assert the truth, and thereby persuade her audience of the truth, then she will pursue that goal with just as much enthusiasm as she would muster if called upon to assert falsehoods. Crucially, however, insofar as she is prompted by her practical interests to assert misleading claims, then she will be aiming for those claims to be convincing, and hence will be trying to mislead her audience. Yet again, then, we find that we are not being presented with a plausible case in which misleading information is deliberately presented as news, and yet there is no desire to actually mislead the target audience.⁹ Accordingly, when bullshit involves the presentation of misleading information as news, then it counts as fake news, but once we understand the details of what is involved in this regard, then this is not something that Gelfert's treatment of fake news can

accommodate.¹⁰

Notice too that if Gelfert were serious about arguing that the intention to mislead is not a property of fake news, then it would be unclear why he claims that satirical magazines like *The Onion* are not fake news. After all, they clearly employ misleading information by design, and may even be misleading and yet their aim is not to mislead. On Gelfert's view, then, what element of fake news do they fail to satisfy? The proper moral to be drawn regarding these satirical magazines is thus that they are not fake news precisely because their aim is not to mislead. Indeed, the reason Gelfert (2018, 106) gives for why these venues are not trading in fake news is that they do not “deliberately” mislead. But of course this entails that an intention to mislead is built into the very notion of fake news.

Another way in which Gelfert's explicit unpacking of his account of fake news is problematic concerns his contention, contrary to my proposal outlined earlier, that fake news is a genuine form of news. For example, he writes that “[...] it must be granted that, just as disinformation is a species of *information*, fake news is [...] a form of *news*.” (Gelfert 2018, 103, *italics in original*) This is a puzzling claim, however, not least because disinformation doesn't seem to be a species of information. Indeed, Gelfert (2018, 104) notes Fred Dretske's (1981, 57) famous—and to my mind devastating—remark in this regard that disinformation is no more a species of information than a decoy duck is a kind of duck.¹¹

So why does Gelfert think it is important to regard fake news as a genuine form of news? His basic reasoning seems to be that anything that reasonably looks like genuine news must be genuine news. Regarding Dretske's point about the decoy duck, for example, he writes that “if something looks like a duck, swims like a duck, and quacks like a duck, it is awfully difficult to recognize that it is not, in fact, a duck.” (Gelfert 2018, 104) But this is a very odd claim for an epistemologist to make in support of the idea that decoy ducks are types of duck. Does it really follow from the purely epistemic fact that it is hard to distinguish decoy ducks from real ducks that the former is genuinely a species of a latter? This seems to confuse an ontological question with an epistemological one.

Gelfert's overarching concern in this regard appears to be the practical one that recipients of the fake news may not be able to tell it apart from the genuine article. As he puts it, “building truth and veracity into the very definitions of ‘information’ and ‘news’—in other words, making them success terms—does little to address the pressing epistemological problem: how to respond to claims presented to us as true by a putative news source, given

that, for all we know, they might (or might not) be fake news.” (Gelfert 2018, 104) But this is confused, on multiple levels. To begin with, our epistemic difficulty in distinguishing X from Y is not a reason at all to treat X as a species of Y, particularly if we have antecedent reason, as we surely do in this case, to not treat it as a species of Y. (Indeed, as we will see below, it is easy to construct cases where genuine news is hard to distinguish from non-news, but where there is no temptation thereby to treat the latter as a species of the former). Moreover, opting on this basis to treat fake news as a genuine species of news doesn’t resolve (or make any more tractable) the ‘pressing epistemological problem’ that Gelfert notes anyway but rather leaves it entirely intact. How does treating fake news as genuine news make it any easier to differentiate between the two? Finally, note that Gelfert is in any case arguing against a straw man here, as no-one has suggested that news should be understood as a success term—i.e., that it must be true in order to qualify as genuine news (even if that is the suggestion regarding information). (Indeed, I will be arguing below that there can be false news).

2. NEWS AND NEWS SOURCES

At this point it will be useful to say something about what news is. I think we need to understand news explicitly in terms of sources of news, in the sense that news is a kind of information that one gets from a news source (at least *qua* news source—the reason for this qualification will become apparent in a moment). What makes a source of information a news source is that it is designed to convey accurate information to others about recent events, where that information is not already widely known. This is why tree rings are a source of information but not a news source—one can extract accurate information from them, but they are not designed for this purpose. In contrast, a newspaper or media outlet is designed to convey information and that’s what makes it a news source. Moreover, it is important to a news source that it is reporting information that is not already widely known, as otherwise merely stating the obvious—repeating old news *ad nauseum*, say—would constitute news. This is also why news is concerned with recent events.¹²

Note that while paradigm cases of news sources like newspapers transmit a great deal of information, there can be genuine news sources that only convey a relatively rudimentary amount of information. For example, the beacon fires that were once used to alert townsfolk

to invasion are clearly news sources, but the accurate information they were designed to transmit was very limited. It is important, however, that a news source has some degree of range in terms of the kind of news that it provides. A one-line message left on the fridge can convey accurate information, and does so by design, but it is hardly a *source* of news (i.e., as opposed to simply being itself a piece of information). Relatedly, it is also important that news sources are relevantly responsive to the world in terms of the information that they convey. A traffic light conveys accurate information about when traffic should proceed and when it should stop, and does so by design, but it isn't a news source as it is determining the facts in question rather than reporting on them. Finally, note that our concern is with the informational output of news sources *qua* news sources. For example, that the website for a particular news source is not well maintained might give us very important information (e.g., about the credibility of the news source), but this is not thereby news.

That news sources are designed to convey accurate information explains why certain apparent sources of information are not genuine news sources. Imagine a computer program that randomly churns out 'reports' on the internet. Even though these reports might be hard to distinguish from genuine news, and hence might be widely treated as news, this clearly isn't a source of news, since it is not even in the business of trying to convey accurate information (even if, from time to time, it succeeds in this respect regardless). (This example also illustrates the earlier point that, *contra* Gelfert, just because something might be hard to distinguish from a genuine news source, as the reports made by this computer might be, it doesn't follow that it is a genuine news source). Indeed, notice how matters change if the computer program was designed to convey accurate information—if, for example, it was constructed in such a way that it provides a digest of reports from other reputable news outlets. Now we would consider it a genuine news source.

Note too that a news source on this view is defined in terms of the motivation to convey accurate information, rather than in outcome-oriented terms, such as whether it generates beliefs, accurate or otherwise, in the audience for that news. There are a number of reasons for this. For one thing, a news source doesn't cease to exist when the audience disappears, much less the intended audience, such that there is no-one whose beliefs could be influenced by this news source. This is why it can make sense to discover the news output of a long-lost civilization (e.g., by unearthing scrolls and such like). But a more important reason is that understanding news sources in terms of what doxastic effect its activities have on its

audience would generate some counterintuitive results. Suppose that one knows that one's audience will believe the opposite of what one tells them. Would that mean that as a news source one should endeavor to systematically feed one's audience opposing information, and hence misinform rather than inform them? In any case, one's audience might be utterly indifferent to the news source, and hence pay it no heed at all in their deliberations, but that wouldn't entail that it wasn't a genuine news source.

One way of putting this last point is that there can be genuine news sources which are nonetheless ineffective in terms of the outcomes they generate (e.g., no-one hears it, or people hear it but ignore it, and so on). The point is that a news source doesn't cease to be genuine just because it is ineffective. Just as we can distinguish between effective and ineffective news sources, we can also distinguish between news sources that are epistemically good and ones that are epistemically bad. Indeed, the goodness of a news source will often be primarily understood epistemically, such that it has a sound epistemic pedigree.¹³ In particular, its goodness resides in it having suitable epistemic properties, such as conveying information that has an adequate epistemic basis. This might mean, for example, that good journalistic practices have been followed—such as checking sources, corroborating information, seeking confirmation/disconfirmation of one's information, and so on—as many of these practices have an epistemic import.¹⁴ Relatedly, it is not a coincidence that a good news source is often reliable, since it is designed to be such. News sources are devoted to conveying accurate information, and so when they are well-designed that's just what they will do. Indeed, they will generate not just reliable information, but information that, when believed by someone on the basis of this news source, will itself have a good epistemic pedigree. It is in this sense that a good news source tends to generate beliefs that will themselves have a positive epistemic standing.

Note that epistemically good news sources can sometimes generate reports that are false or misleading, including reports that are intentionally so (as when an errant reporter ignores all protocols and invents a story in order to meet a deadline). Since the epistemic goodness of the source relates to its structural epistemic properties and its general reliability, this is compatible with it sometimes not successfully meeting its epistemic goals. This is why it does not follow on this way of thinking about news that it is a success term—the reports made by a genuine news source can be false, and even intentionally so, and thus there can be in this sense false news (which, note, is not thereby fake news, but rather genuine news, albeit

of an epistemically problematic variety).

Some news sources are not epistemically good in this sense, however. For example, a school newspaper might not incorporate any of the protocols that are characteristic of epistemically good news sources. It would follow that this news source would tend to deliver reports that are unreliable. But notice that a news source that is bad in this epistemic sense can still be a genuine news source. This is because it is still aiming to convey accurate information of the relevant kind, it is just that it isn't very good at doing so. Remember that we have defined news sources in terms of manifesting a certain kind of intentional activity, and hence one can instantiate this notion well or poorly. That one is not very good at archery, and so regularly misses the target, doesn't mean that what one is doing when one fires one's arrows at the target isn't archery. Of course, if one is consistently hopeless at achieving the relevant goal of that activity, even when to do so would be relatively straightforward, then that starts to call into question whether one is genuinely striving to achieve it in the first place. But the point remains that if we understand what makes something a news source, then it becomes clear that there can be epistemically bad news sources as much as there can be epistemically good news sources, just as there can be ineffective as well as effective news sources.¹⁵

3. FAKE NEWS *VERSUS* NEWS

As should be plain from the foregoing, fake news is not to be understood as a kind of news, albeit of a problematic variety, as Gelfert proposes. It is, rather, not genuine news at all, since it doesn't meet the requirement of being aimed at conveying accurate information of the relevant kind. Instead, as the 'fake' tag indicates, it is masquerading as real news in order to spread misinformation. Note too that fake news is different from other forms of non-news, such as the rings in the tree, the computer program randomly churning out reports, or the satirical magazine (even supposing that these sources deliver new information about recent events, which of course they usually don't). In the first case, there is no misinformation, let alone misinformation by design. In the second and third cases, even if there is systemic misinformation, there is no intent to misinform.

Moreover, notice that just as fake news is not a form of genuine news, so *a fortiori* it is not a form of epistemically bad news either. That is, it is not that fake news is genuine news

that has epistemically problematic properties, which is the case with epistemically bad news. Even though both phenomena are epistemically problematic, they are distinct in that epistemically bad news is still a variety of news, in virtue of being in the business of trying to convey accurate information (albeit poorly). This point is particularly important to adjudicating whether some cases are fake news or merely epistemically problematic news. For example, we can imagine an agent who is in the grip of various conspiracy theories and hence regularly churning out false stories as news. It is crucial here, however, whether in doing so they are trying to convey accurate information. If they are, then their output isn't fake news, but rather genuine news that is simply epistemically deficient. It is still epistemically problematic, but the problem it poses is distinct from that raised by fake news.¹⁶

Another problematic feature of Gelfert's account is that although he doesn't demand that fake news involves an intent to mislead, he nonetheless advances a very strong thesis regarding the extent to which fake news must actually mislead. He writes that fake news "must *in fact* mislead a relevant audience" (Gelfert 2018, 103, *italics in original*), that it is important to fake news that "sufficiently large numbers of people are in fact taken in by it" (Gelfert 2018, 105), and that fake news must be "*objectively likely* to mislead its target audience." (Gelfert 2018, 108, *italics in original*)¹⁷ These are demanding requirements, and they do not stand up to closer scrutiny. The crux of the matter is that just as we need a distinction between effective news and ineffective news, such that the latter can nonetheless be genuinely news, so we need to distinguish between effective and ineffective fake news, where the latter is nonetheless a *bona fide* form of fake news. In particular, we need a way of distinguishing between kinds of fake news that are constructed such that they are effective at attaining their epistemic objective of misleading people as opposed to kinds of fake news that are constructed such that they are ineffective at attaining these epistemic objectives.

Consider, for example, a website that is designed to spread misinformation by embedding false stories within accurate ones, perhaps for propaganda purposes. This would be a paradigmatic case of fake news. Nonetheless, this can be done in an effective way, or in a lousy way. In the former case, this would involve carefully curating the false reports within the accurate information in order to make it more compelling, phrasing the false reports judiciously to ensure that they look as plausible as possible, offering ready-made responses to natural objections that the reader might have to the content of these reports, and so on. The upshot would be to provide an epistemic setting that makes the fake news more effective, in

virtue of it being harder to distinguish from genuine news. In the latter case, in contrast, one could imagine the very same website conducting its business in a clumsy way, such that it will be relatively easy for an informed person to spot that the news is fake. In particular, the epistemic stage-setting will be lacking, and hence the alarm bells will sound for most people that this website is not to be trusted.

The point of the matter is that poorly executed fake news is still fake news. In particular, the fact that the lousy fake news website might not in fact mislead anyone, and certainly wouldn't be such that it is 'objectively likely' that it would mislead its intended audience, doesn't entail that it isn't fake news. *Modulo* our earlier point about decoy ducks and real ducks, just as an excellent forgery of a Picasso painting does not become the genuine article just because it is so good as to be hard to distinguish from the real thing, so it doesn't follow that a lousy forgery of a Picasso painting is any less of a forgery for the fact that it is easily spotted. In short, ineffective fake news is still fake news.

More generally, just as we should avoid restrictive accounts of fake news that demand that fake news always be effective, so we should similarly reject proposals that build specific goals or purposes into the account of fake news, such that fake news is identified with a particular form of fake news. Consider, for example, this account of fake news offered by Scott Aikin and Robert Talisse:

“Fake news characterizes the activities of institutions that *pose* as journalistic which by design feed and codify the *antecedent* biases of a pre-selected audience by exploiting their vulnerabilities (cognitive and otherwise), all with a view towards facilitating some decidedly political objective.” (Aikin & Talisse 2018, *italics in original*)

While the proposal captures the point that fake news is not a kind of genuine news (but merely masquerades as such), it offers a much too narrow conception of what fake news involves. On this account, in order for someone to create fake news they must already have a sophisticated conception of cognitive biases and how to effectively exploit them. They must also have political objectives in doing this. But neither seems essential to fake news. On the former front, while it might be a feature of (a certain form of) *effective* fake news that it is designed to exploit cognitive biases in this way, we have already noted that there can *ineffective* forms of fake news which are no less genuine varieties. Moreover, while many forms of fake news surely do have a political agenda, it's hard to see why this should be an essential feature, as fake news could serve many purposes (including merely being, say, for the amusement of the one propagating it).

Putting all these points together, let's review where we are with our account of fake news. We've argued that it has the following properties. First, unlike genuine news (which is aimed at conveying accurate information), fake news deliberately conveys misleading information. This doesn't mean that fake news necessarily involves the presentation of falsehoods, given that even the literal truth can be misleading (indeed, as we've noted, sometimes this is the most effective way to mislead). Nonetheless, second, fake news is presented as news, just as a forgery of a painting is presented as the real thing. Relatedly, third, fake news involves an intent to mislead, which is why, for example, satirical news magazines are not fake news (even if they in fact do mislead). That said, fourth, fake news need not in fact mislead. While effective fake news typically misleads, ineffective fake news may not be successful on this front, but it is no less a form of fake news as a result (ineffective fake news is still fake news). Finally, fifth, fake news is not itself a type of news, even though it may be hard to distinguish between the two (any more than an excellent forgery is thereby the genuine article). In particular, fake news is to be distinguished from a genuine form of news that has a poor epistemic pedigree.

4. MANAGING FAKE NEWS: INDIVIDUAL

One question we might ask at this juncture is what hangs on our thinking about fake news in just this fashion. I think we can recognize the importance of understanding fake news correctly once we appreciate the kind of expertise needed to spot it. In particular, since fake news is to be distinguished from genuine news, even genuine news that has a poor epistemic pedigree, what is required to spot fake news won't be quite the same as what is required to assess the epistemic credentials of genuine news (though there will be quite a lot of overlap, as we will see).

This ought not to be surprising. Consider someone who works in a treasury whose job is to ensure the quality control of the currency that is being produced. This will require very specific skills to spot the kinds of considerations that are relevant to currency quality, where this might include such things as printing issues, degradation of the (e.g.,) polymer on which the note is printed, errors in what is printed, and so on. Naturally, someone working in this department will also be interested in how the notes were produced, particularly where the

quality is lacking, so that they can ensure that future quality is preserved.

Someone with these skills would likely be well-placed to spot a counterfeit note, not least due to their close familiarity with genuine currency, but it remains that the skills involved in spotting counterfeit currency are not quite the same. After all, our treasury currency controller is not dealing with fakes, but rather trying to spot genuine currency that is merely deficient in some way. In particular, there is no element of deceit in play in this quality control process, but merely a sifting of the wheat from the chaff.

In contrast, we can imagine someone who works at the same treasury whose specific job it is to spot counterfeit currency. While this person will have similar expertise to her colleague in quality control, there will clearly be differences to how she goes about her role. The counterfeit currency might not have the kinds of errors or deficiencies that are found in low quality currency; indeed, one would expect the very best counterfeit currency to be very similar to genuine currency, at least superficially. Our expert in the counterfeit currency department will thus be on the look-out for very specific indications that a counterfeit is in play (an unexpected barcode, a watermark that is just a little too large or not quite in the right place, and so on). Moreover, those working in the counterfeit department of the treasury won't be looking just at the notes, but will also be interested in other information regarding the provenance of the notes—indeed, one would expect that it is information of just this kind that has led to these notes being brought to this department for inspection.

What goes for currency goes for news, or so I claim. While there will be obvious overlap in the kind of skills and knowledge that is required to distinguish epistemically good genuine news from epistemically deficient genuine news, and fake news from genuine news, there is nonetheless a distinct set of skills and knowledge in play in each case. We need to attend to this fact if we wish to train people to spot fake news.

We can capture what is at issue here by comparing what is required to evaluate the epistemic credentials of testimony in a context in which testifiers are generally honest (but might not thereby always have a good epistemic basis for what they assert, much less always assert the truth), and what is required to evaluate the epistemic credentials of testimony in a context where there are bad actors in play who have an intent to deceive (note that in keeping with our account of fake news, this might not involve lying as such). This would be roughly comparable to our quality control person and our counterfeit spotter at the treasury described above, except now the 'currency' under evaluation is testimony. There is one interesting

difference between the two cases, however. The reason why people forge currency is very straightforward, as they are obviously seeking financial gain. The reasons why people use testimony to deceive, however, just as the reasons why people propagate fake news, are numerous, and might not be at all straightforward. This is important, since it gives the person evaluating the testimony more to go on when it comes to undertaking their evaluation.

In a testimonial context where one is assured that the testifiers are sincere, one needs only to consider such factors as the reliability of the testifier and the credibility of what they are claiming. In contrast, in testimonial contexts where one lacks such an assurance, then one has to in addition consider the motivations these testifiers might have for asserting what they do. In particular, one needs to consider whether the testifiers have a motivation to mislead by offering this particular testimony. The epistemic burdens on those assessing the testimony are thus higher, and what is required to weed out good testimony from bad is correspondingly more demanding.

What goes here for testimony also applies to the task of identifying fake news. This should be unsurprising, given that being a recipient of (putative) news is itself a kind of testimonial context (albeit one that doesn't typically involve any direct interaction between the testifier and the recipient of the testimony). In particular, if one is on the lookout for fake news, then one needs to take into account the specific motivations someone might have for propagating misleading information. This is different from the case where one is merely distinguishing between genuine news that enjoys a good epistemic pedigree and genuine news that doesn't. In that case, one's concern is entirely with the epistemic source of the news (e.g., does it arise from a reliable news source?), and with the epistemic standing of the news itself (e.g., is it credible, given what else one knows?). Since the sincerity of those offering this news is not in question, one doesn't need to worry about their motivations. When one is in addition on the look-out for fake news, however, then motivations become very salient. In particular, one has to further consider such factors as whether this is 'news' that someone might have a motivation to put out even though it is misleading. For example, where the 'news' benefits a particular political party during an election campaign, then one should be willing to expose it to additional scrutiny, since this fact raises the likelihood that it might be fake news.

We saw earlier that there will be considerable overlap in the expertise required to, on the one hand, differentiate good quality currency from poor quality currency, and, on the other hand, differentiate genuine currency from counterfeit currency. The same will be true in

the case of genuine and fake news, especially when it comes to epistemically deficient genuine news and ineffective fake news. For example, a genuine news source might, through error, end up reporting something clearly false, and which an alert recipient of this news can easily spot is false. Similarly, obvious falsehoods might be propagated as fake news, and easily spotted in much the same way.

The interesting contrast, however, will involve the hard cases of differentiating genuine news from effective fake news. This is where the special expertise relevant to spotting fake news—just as in the case of spotting counterfeit currency—becomes important. Effective fake news, after all, will be by its nature hard to distinguish from the genuine article. In particular, it might not be obvious either that the source of the putative news is epistemically problematic or that what it claims isn't credible. This is where attending to possible motives for putting forward these claims, and relatedly being aware of the wider context of the putative news (political, social, commercial etc.), becomes important.

We can see some of these points in action by considering some recent empirical work on how people are taken in by fake news. As researchers have noted, the problem is that recipients of the fake news tend to evaluate the putative news source 'vertically' rather than 'horizontally', where this means that they look for corroborating evidence (or otherwise) for what they are reading from *within* the article itself rather than verifying the putative news source by appeal to independent, and already verified, sources of news.¹⁸ Accordingly, they find themselves becoming more convinced of the fake news the more they read of it, when in fact they ought to be far more suspicious of what they are reading (and would be, had they examined relevant independent, and already corroborated, sources of news).

The mistake our recipients of fake news are making is to fail to realize that they are not in a testimonial context where they can reasonably take it as given that the testimony they are receiving is honest and sincere. In particular, they are failing to realize that they cannot assume that what seems like genuine news is genuine news. Were they to be in such an epistemically friendly testimonial context, then a vertical epistemic evaluation would be far more plausible as a system of epistemic appraisal, just as one might epistemically evaluate someone's testimony in such a testimonial context by considering how credible it is that the claims they are making are true. (Though as we noted above, even then one should also be willing to consider the epistemic standing of the source of the testimony, as even honest and sincere testifiers can be unreliable, particularly about certain subject matter; being responsive to the

credibility of what is asserted would be a natural first step, however). In testimonial contexts where one cannot take it as given that the testimonial actors are honest and sincere, however, then a further layer of scrutiny is required, and that's where the need for a horizontal epistemic evaluation becomes pressing, for how else is one to satisfy oneself that this apparently genuine news is the real deal, and not its fake counterpart? The challenge is thus to have the cognitive skills to know when it is appropriate to adopt the additional levels of scrutiny, and a big part of that, as we noted above, is being sensitive to the possible motives that someone might have to make the target 'news' claims in play.

With all these points in mind, what kinds of cognitive traits should we promote in subjects if we want them to be able to spot fake news? I think the answer lies in the *intellectual virtues*, where by this I mean those distinctive admirable character traits that are guided by a love of the truth, and which are essential components of a life of flourishing, such as intellectual humility, intellectual conscientiousness, honesty, and so on.¹⁹ There is no infallible guide to spotting fake news, but as fallible guides go I believe the intellectual virtues have the most promise.²⁰

To begin with, if we are to counter the threats to truthfulness posed by fake news, then it is essential that we foster character traits that involve a love and desire for the truth, as lack of concern for the truth is precisely what is generating the problem in hand. It is thus significant that the intellectual virtues essentially incorporate such a desire for truth as a motivational component. Relatedly, it is also important to the intellectual virtues that these are not cognitive traits that are innate, or of a kind that can be easily acquired and maintained. The intellectual virtues require *cultivation*, and it is the process of such cultivation, given the truth-directed nature of these character traits, that keeps the value of truth and its possession as a target in mind. The proponent of fake news might not care about the truth, but the recipient of it needs to have this epistemic good in her sights.

Even more importantly for our purposes, however, a key facet of the intellectual virtues is the manner in which they are highly context-sensitive in their application. That is, it is part of the very nature of an intellectual virtue that it involves a sensitivity on the part of the subject to relevant circumstances. Indeed, this much follows from the fact that the manifestation of the intellectual virtues (in common with the virtues more generally) lies on a 'golden mean' between two vices, one of excess and one of deficiency. So, for example, the intellectual virtue of being intellectually conscientious lies between the vice of excess of

pedantically attending to every detail, no matter how irrelevant, and the vice of deficiency of being intellectually unconscientious (e.g., being completely unconcerned about the relevant evidence). To have an intellectual virtue is thus to have the good judgement to know when to manifest this cognitive trait, such that one is neither deficient nor excessive in the relevant respects.

This is significant when it comes to fake news—just as it is relevant to the epistemic assessment of testimony more generally—in that it means that the intellectually virtuous subject will be sensitive to those factors that demand further inquiry. Whereas it might often be appropriate to evaluate putative news sources on the assumption that they are genuine (in that subjecting them to additional scrutiny would be excessive), and hence make broadly vertical epistemic assessments of them, the intellectually virtuous subject will also be sensitive to considerations that would demand that one should raise one's level of scrutiny. In particular, the intellectually virtuous subject will be responsive to those factors that would entail that this assumption about putative news sources being genuine should be dropped (such as the motivational considerations that underlie fake news noted above), and hence that a horizontal epistemic assessment would be appropriate instead.²¹

It is also relevant in this regard that the intellectual virtues are by their nature integrated with one another, and more generally with the other (non-intellectual) virtues, such as the moral virtues, with both sets of virtues governed by an overarching virtue of practical wisdom (*phronesis*). This feature of intellectual virtues, coupled with the fact that they are manifested on a golden mean between two corresponding vices of excess and deficiency, ensures that the virtuous agent responds to the threat posed by fake news in a measured manner. In particular, there is a contemporary tendency to deal with this threat by becoming overly sceptical about news sources more generally, and in the process dismissing much that has epistemic merit. Indeed, in the more extreme case this can lead people to be dismissive even of scientific reporting from reputable sources. In virtue-theoretic terms, this is to succumb to a vice of excess, in contrast to a more proportionate, and intellectually virtuous, questioning of sources in light of the presence of fake news in one's epistemic environment.²²

Moreover, notice that such widespread (and non-virtuous) scepticism is an intellectual strategy that has become detached from one's wider practical and moral goals, something that the integrated nature of the virtues is designed to prevent. In particular, to doubt in this fashion is to cut oneself off from epistemic goods and thereby lose the practical and moral

advantages that the possession of these goods can provide. In the case of one's practical goals, the import of such epistemic goods ought to be straightforward. How is one to achieve one's practical goals if one is suspicious of news sources more generally, given how such sources can be vital to informing one's decision-making? I think the relevance of such epistemic goods to one's moral goals is also relatively clear. My claim here, in common with virtue theoretic accounts more generally, is that knowledge is required for one to be morally virtuous, in that relevant knowledge (e.g., about one's circumstances) is needed to guide the manifestation of the moral virtues. It follows that to deprive oneself of knowledge via such widespread scepticism can prevent one from manifesting these admirable character traits. Given the importance of the virtues to a life of flourishing, to undermine one's manifestation of the moral virtues in this way is effectively a form of self-harm. In any case, possession of the virtues, including the intellectual virtues, acts as an important barrier towards taking such a route.

5. MANAGING FAKE NEWS: STRUCTURAL

In the last section we focused on what consequences our account of fake news has for the individual who needs to be able to spot it. But clearly it cannot only be exclusively the job of the individual to manage fake news; there also needs to be a structural response to this wider social problem. Aside from anything else, even the most intellectually virtuous individual might lack the capacity to differentiate genuine news from its counterfeit counterpart; as we noted, while the intellectual virtues are plausibly the best route to making such a differentiation (at the individual level anyway), they are far from being infallible guides. This is why we need to assist the individual by in addition seeking structural reforms to our epistemic environments that help us to identify fake news.

Our account of fake news is helpful in this respect. This is because one of the challenges facing those who seek structural responses to the problem of fake news is the danger that this leads to putting undue constraints on a free press, something that is widely held to be a cornerstone of a well-functioning democratic society. But notice that this concern betrays a way of thinking about fake news that regards it as being continuous with genuine news, along the lines that we saw Gelfert setting out above. It is only on this supposition that

any structural constraints on fake news would *thereby* be an infringement on genuine news. In contrast, insofar as we treat fake news as distinct from genuine news, as we have been urging, then the conceptual space is cleared to allow us to regulate the former without thereby undermining the latter.

Of course, there will undoubtedly be practical hurdles to clear in this regard, as any policy of containing fake news will potentially have a bearing on the practices of those putting forward genuine news, especially since it can be hard, as we have seen, to differentiate effective forms of fake news from the real thing. But such practical hurdles are far from unsurmountable, and at most all they entail is that we should proceed with caution. So, for example, it might be necessary to set the bar for fake news, of a kind that is to be regulated at a structural level anyway, quite high, such it is only clear-cut cases of fake news that get curtailed.

It is also obviously relevant here what this ‘curtailment’ involves and who is doing it. In the former case, if all that is happening is that the ‘news’ comes with a warning to the recipient that it might be fake news, then that hardly seems to be infringing on one’s democratic freedoms at all, particularly *modulo* the previous point that we are only picking out clear-cut cases of the phenomenon. In contrast, if ‘curtailment’ means actual removal of the fake news, then that will obviously be more problematic in this regard. This point dovetails with the issue of who is doing the censoring. Where this is itself democratic institutions that are open to challenge from its citizens, then again it will be harder to make the case that this is an infringement of one’s democratic freedoms. In contrast, if this censorship is done by non-governmental bodies, such as corporations (e.g., Facebook, Twitter, and so on), especially if this is not subject to outside scrutiny or governmental regulation, then I think we will rightly feel somewhat nervous about how this might lead to a wider constraint of genuine journalistic reporting.

It is not my goal here to navigate through these difficult practical waters, but merely to note that it makes a crucial difference to these debates whether one holds that fake news is itself a kind of genuine news, albeit of a problematic epistemic kind. Once one departs from this supposition—and I have argued that we should do so—then it becomes possible to identify ways of dealing with fake news that need not thereby impose a restriction on genuine news (including genuine news that has a problematic epistemic pedigree). Moreover, although the task of differentiating genuine news from fake news at the individual level may be arduous

(even if one grants the subject the intellectual virtues that I have described above), the combination of the intellectually virtuous subject working in concert with a structural approach to fake news should make the task of identifying fake news far more tractable.²³

REFERENCES

- Aikin, S. F., & Talisse, R. (2018). 'On 'Fake News'', *3 Quarks Daily*, <https://www.3quarksdaily.com/3quarksdaily/2018/05/on-fake-news.html>.
- Battaly, H. (2014). 'Intellectual Virtues', (ed.) S. van Hooft, *Handbook of Virtue Ethics*, 177-87, London: Acumen.
- Bachr, J. (2011). *The Inquiring Mind: On Intellectual Virtues and Virtue Epistemology*, Oxford: Oxford University Press.
- (2013). 'Educating for Intellectual Virtues: From Theory to Practice', *Journal of Philosophy of Education* 47, 248-62.
- (2019). 'Intellectual Virtues, Critical Thinking, and the Aims of Education,' *Routledge Handbook of Social Epistemology*, (eds.) M. Fricker, P. Graham, D. Henderson & N. J. L. L. Pedersen, ch. 43, London: Routledge.
- Carter, J. A., Kotzee, B., & Siegel, H. (2019). 'Educating for Intellectual Virtue: A Critique from Action Guidance', *Episteme*, [Online First, DOI: <https://doi.org/10.1017/epi.2019.10>].
- Bernecker, S. (Forthcoming). 'An Epistemic Defense of News Abstinence', *Epistemology of Fake News*, (eds.) S. Bernecker, A. Flowerre & T. Grundman, Oxford: Oxford University Press.
- Cassam, Q. (2018). 'Epistemic Insouciance', *Journal of Philosophical Research* 43, 1–20.
- Dretske, F. (1981). *Knowledge and the Flow of Information*, Cambridge, MA: MIT Press.
- Fallis, D., & Mathiesen, K. (2019). 'Fake News is Counterfeit News', *Inquiry*, [Online First, DOI: <https://doi.org/10.1080/0020174X.2019.1688179>].
- Frankfurt, H. G. (2005). *On Bullshit*, Princeton: Princeton University Press.
- Gelfert, A. (2018). 'Fake News: A Definition', *Informal Logic* 38, 84-117.
- Grundmann, T. (2019). 'Fake News: The Case for a Consumer-Oriented Explication', *manuscript*.
- Habgood-Coote, J. (2018). 'Stop Talking about Fake News!', *Inquiry*, [Online First, DOI: <https://doi.org/10.1080/0020174X.2018.1508363>].
- Jacquette, D. (2010). 'Journalism Ethics as Truth-Telling in the Public Interest', *Routledge Companion to News and Journalism*, (ed.) S. Allen, 213-22, London: Routledge.
- Levy, N. (2017). 'The Bad News About Fake News', *Social Epistemology Review and Reply Collective* 6, 20-36.
- MacKenzie, A., & Bhatt, I. (2018). 'Lies, Bullshit and Fake News: Some Epistemological Concerns', *Postdigital Science and Education*, [Online First, DOI: <https://doi.org/10.1007/s42438-018-0025-4>].
- McGrew, S. (2019). 'Learning to Evaluate: An Intervention in Civic Online Reasoning', *Computers & Education*, [Online First, DOI: <https://10.1016/j.compedu.2019.103711>].
- McGrew, S., Ortega, T., Breakstone, J., & Wineburg, S. (2017). 'The Challenge that's Bigger than Fake News: Civic Reasoning in a Social Media Environment', *American Educator* 41, 4-11.
- McGrew, S., Breakstone, J., Ortega, T., Smith, M., & Wineburg, S. (2018). 'Can Students Evaluate Online Sources? Learning from Assessments of Civic Online Reasoning', *Theory & Research in Social Education* 46, 165-193.
- (2019). 'Improving University Students' Web Savvy: An Intervention Study', *British Journal of Educational Psychology* 89, 485-500.
- Lazer, D. M. J., Baum, M. A., Benkler, Y., Berinsky, A. J., Greenhill, K. M., Menczer, F., Metzger, M. J., Nyhan, B., Pennycook, G., Rothschild, D., Schudson, M., Sloman, S. A., Sunstein, C. R., Thorson, E. A., Watts, D. J., & Zittrain, J. L. (2018). 'The Science of Fake News', *Science* 359, 1094-96.
- Mukerji, N. (2018). 'What is Fake News?', *Ergo* 5, 923-46.
- Pepp, J., Michaelson, E., & Sterken, R. K. (2019). 'What's New About Fake News?', *Journal of Ethics*

- and Social Philosophy* 16 [Online First: <https://doi.org/10.26556/jesp.v16i2.629>].
- Pritchard, D. H. (2013). 'Epistemic Virtue and the Epistemology of Education', *Journal of Philosophy of Education* 47, 236-47.
- (2016). 'Intellectual Virtue, Extended Cognition, and the Epistemology of Education', *Intellectual Virtues and Education: Essays in Applied Virtue Epistemology*, (ed.) J. Baehr, 113-27, London: Routledge.
- (2018a). 'Disagreement, of Belief and Otherwise', *Voicing Dissent: The Ethics and Epistemology of Making Disagreement Public*, (ed.) C. Johnson, 22-39, London, Routledge.
- (2018b). 'Neuromedia and the Epistemology of Education', *Metaphilosophy* 49, 328-49.
- (2019). *Scepticism: A Very Short Introduction*, Oxford, Oxford University Press.
- Rini, R. (2017). 'Fake News and Partisan Epistemology', *Kennedy Institute of Ethics Journal* 27, 43-64.
- Rose, J. (2019). 'To Believe or Not to Believe: An Epistemic Exploration of Fake News, Truth, and the Limits of Knowing', *Postdigital Science and Education*, [Online First, DOI: <https://doi.org/10.1007/s42438-019-00068-5>].
- Siegel, H. (2016). 'Critical Thinking and the Intellectual Virtues', *Intellectual Virtues and Education: Essays in Applied Virtue Epistemology*, *Intellectual Virtues and Education: Essays in Applied Virtue Epistemology*, (ed.) J. Baehr, 95-112, London: Routledge.
- Siegel, H. (2018). *Education's Epistemology: Rationality, Diversity, and Critical Thinking*, (Oxford: Oxford University Press).
- Twardowski, K. (1979). 'Issues in the Logic of Adjectives', *Semiotics in Poland 1984–1969*, (ed.) J. Pelc, 289-30, Dordrecht, Holland: Springer.
- Wineburg, S., & McGrew, S. (2017). 'Lateral Reading: Reading Less and Learning More When Evaluating Digital Information', *Stanford History Education Group Working Paper No. 2017-A1*. [Available at SSRN: <https://ssrn.com/abstract=3048994>]
- Zagzebski, L. T. (1996). *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*, Cambridge: Cambridge University Press.

NOTES

¹ This is, of course, how Donald Trump famously uses this terminology.

² For a defence of the very different philosophical stance that ‘fake news’ should not be given a philosophical treatment, on account of how it doesn’t pick out a specific phenomenon, see Habgood-Coote (2018). If the philosophical account that I will be offering of fake news is credible, then it is thereby a (partial) response to this particular critical line.

³ For some closely related philosophical discussions of fake news, see Levy (2017), Rini (2017), and Rose (2019).

⁴ This is one of the problems that faces the account of fake news offered by Lazer *et al* (2018), as their proposal explicitly defines fake news in terms of ‘fabricated information’. But as Aikin & Talisse (2018) point out, one can in fact use accurate information to mislead, such as by presenting it in a way that triggers a cognitive bias.

⁵ At least, false claims are characteristically misleading, if not universally so. For example, I am here setting aside such unusual cases as when one knows that one is dealing with someone who consistently lies, in which case their false claims might well be reliable ways to discern the truth, and hence not misleading at all.

⁶ This is a point that Rini (2017) seems to miss in her discussion of fake news:

“[...] I said that fake news requires intentional deception, but this may be too strong. Deception is not always the primary goal of fake news. Often the motive is financial rather than epistemic.” (Rini 2017, §2)

Notice how the idea of fake news involving an intent to mislead is converted into this being the ‘primary goal’ of the fake news. But it is entirely consistent with the idea that an intention to mislead is built into one’s account of fake news that the purveyors of fake news also have other motives in play (and indeed that those other motives play a more primary role in their activities). Note too that Rini’s account of fake news is also problematic in that she insists that it involves ‘more than mere lying’. As we have noted, however, while fake news is designed to mislead, one does not need to lie (i.e., assert a falsehood) in order to do that.

⁷ See also Cassam’s (2018) discussion of the related cognitive trait of epistemic insouciance.

⁸ I don’t think the case of bullshit is unique in this regard, as there are a number of epistemologically interesting cases of assertion where what might initially look like the expression of belief is in fact nothing of the kind, at least in any robust sense of belief of a kind that would be of interest to epistemologists (e.g., which is a constituent part of rationally grounded knowledge). For further discussion of this point, and some of its epistemological ramifications, see Pritchard (2018a).

⁹ I think this point also counts against one of the reasons that Pepp, Michaelson & Sterken (2019) offer for thinking that fake news needn’t involve an intention to mislead:

“Consider an organization like the *National Enquirer*. Suppose that the motivation of the proprietors of this organization is to maximize profit, and they incentivize their employees in order to carry out this maximization. It is not all that hard to imagine that each employee of the *National Enquirer* might act either (a) to maximize profit, or (b) to maximize personal gain, with no one intentionally putting forward content with the intention to deceive or mislead anyone else.” (Pepp, Michaelson & Sterken 2019, 74)

But how on earth would these actors ensure that this strategy is successful *without* an intention to mislead? After all, they will want their output to be read and believed, and they know that this will result in their readers being misled. That the readers are misled thus doesn’t seem to be a merely predictable (but unintended) consequences of their actions, but one that is actively sought. Remember too that it is quite consistent with the claim that fake news involves an intention to mislead that it also involves other motivations, and indeed that those other motivations are more central to the activity.

¹⁰ See Mukerji (2018) for an account of fake news as essentially bullshit presented as news. I think Mukerji’s analysis overlooks the issue raised here, however. This is because he seems to assume that it follows from Frankfurt’s account of bullshit that the bullshitter, by being indifferent to the truth, is thereby also indifferent to whether they are believed. Consider this passage:

“Unlike a liar, however, who seeks to convince us that a given statement he believes to be false is true, a bullshitter does not care about whether his utterances are true.” (Mukerji 2018, §2)

But this is the wrong contrast to draw. Unlike the liar, it isn’t essential to the bullshitter’s claims that he is seeking to convince us that *what he believes is false* is true, but he *is* trying to convince us of his claims (it’s just that it doesn’t matter to him whether they are true or false). Relatedly, bullshit *qua* fake news also involves an intention to mislead. Since he misses this point, Mukerji falsely supposes that all that is required on this score on the part of the bullshitter is the intention to mislead the audience about his true motives in making his claims (which of course doesn’t entail any narrow intention to mislead with regard to the content of the fake news itself). Note too that Mukerji’s account is also unable to accommodate the phenomenon of fake news that involves straightforward lying, as he himself concedes (Mukerji 2019, §5).

¹¹ To use the terminology offered by Twardowski (1979), an adjective like ‘decoy’ when applied to ‘duck’ is an ‘eliminating’ adjective, in that it is suggesting that the target item (the duck) is not present. (Another example of an expression using such an adjective is ‘false friend’). This is in contrast to ‘determining’ adjectives (like ‘red’ in ‘red coat’) or

‘confirming’ adjectives (like ‘actual’ in ‘actual fact’) which don’t have this implication. I am grateful to an anonymous referee for alerting me to this point.

¹² Of course, the news could be ultimately concerned with events in the distant past, as when one makes a discovery regarding the ancient world and this becomes news. But it is the discovery which makes it news, and that is a recent phenomenon. Relatedly, news can be regarding the future, as when a newspaper reports scientific claims regarding the anticipated effects of global warming. But again what makes this news are the scientific claims, and they are recent. Finally, notice that strictly speaking on this proposal we should be indexing news to a specific time, in that what was previously genuine news will likely not be news in the future (at least assuming that nothing relevant changes in the interim). So old news is no longer a genuine form of news, but it once was (unlike other forms of information which were never forms of news). Since this complication doesn’t concern us here, I will be setting it to one side in what follows.

¹³ As should be clear, when I talk of ‘good’ news in this context, I am specifically concerned with the epistemic properties of the news source, and not whether the news it generates is to be otherwise welcomed. The same goes, *mutatis mutandis*, with ‘bad’ news. In what follows I will take this point as granted.

¹⁴ Journalistic practices are not all aimed at enhancing the epistemic pedigree of the news source, of course, as they can also serve other ends, such as principles of fairness. For further discussion of journalistic practices from an epistemic point of view, see Jacquette (2010).

¹⁵ I think this is a point that is overlooked in some of the literature on fake news. For example, Fallis & Mathieson (2019) and Pepp, Michaelson & Sterken (2019) both understand news in a way that incorporates good epistemic procedures into the news itself (via appeal to appropriate journalistic practices). But this is to conflate news with a good epistemic pedigree with news *simpliciter*, and thereby to fail capture the fact that there can be news that has a poor epistemic pedigree but is no less news as a result.

¹⁶ For a contrasting view, see Pepp, Michaelson & Sterken (2019) who claim that such a case would be an instance of fake news. As I’ve just explained, I think this is to confuse fake news with genuine news that is epistemically problematic. As noted in endnote 15, part of the reason for our divergence is that Pepp, Michaelson & Sterken (2019) build far more into a news source than I do, as they identify news sources with those that manifest good journalistic practices. But I think this is to equate news sources with good news sources (where the ‘goodness’ in play is primarily epistemic).

¹⁷ See also Grundmann (2019) and Bernecker (*forthcoming*), who both define fake news in terms of the effects it has on its audience, such that fake news needs to be effective in misleading its intended audience.

¹⁸ See especially Wineburg & McGrew (2017). This paper is part of a rich body of recent work on how people evaluate digital news sources that has been produced by the Stanford History Education Group. For some other representative publications arising out of this research, see McGrew *et al* (2017; 2018; 2019) and McGrew (2019). I am grateful to Gabe Avakian Orona for alerting me to this work.

¹⁹ I take a broadly Aristotelian line on the intellectual virtues, as my subsequent discussion will reveal. For a key contemporary defence of such an account of the intellectual virtues, see Zagzebski (1996). See also Baehr (2011). For a more general overview of the contemporary literature on intellectual virtue, see Battaly (2014).

²⁰ At least at the individual level anyway—the reason for this qualification will soon become apparent. Incidentally, elsewhere I have argued at length for the claim that the overarching epistemic goal of education is the cultivation of the intellectual virtues in just this sense—see, especially, Pritchard (2013; 2016; 2018*b*). See also Baehr (2013) for a proposal in a similar vein. If that’s right, and if it is also correct that the intellectual virtues offer the best way for us to identify fake news (at least at the individual level), then this educational goal will also serve this more specific purpose. For a recent critical discussion of this kind of account of the epistemology of education, see Siegel (2016; 2018) and Carter, Kotzee & Siegel (2019). See Baehr (2019) for a response to this critical line.

²¹ This is one reason why I am sceptical about educational attempts to deal with fake news that simply focus on providing students with various kinds of technical expertise, such as knowledge of statistics and how they can be misused. While such technical expertise is undoubtedly useful in this regard, the point is that one also needs to cultivate the good judgement in the students to employ this expertise appropriately, and for that the intellectual virtues are required. This recent national schools project in Finland, as reported in *The Guardian* newspaper, is interesting in this regard:

□ <https://www.theguardian.com/world/2020/jan/28/fact-from-fiction-finlands-new-lessons-in-combating-fake-news>

At least on the surface, the project seems to be primarily about the development of technical expertise, but if one looks closer at how this project is described it begins to look far more about the more general cultivation of intellectual character, and thus intellectual virtue.

²² I explore this theme in more detail in Pritchard (2019).

²³ Thanks to Sven Bernecker, Michel Croce, Thomas Grundman, and Gabe Avakian Orona. Material from this paper was presented at the ‘Archaeology in a Post-Truth World’ conference, held at UC Irvine in February 2020. I am also grateful to two non-anonymous reviewers for this volume for their detailed comments on an earlier version of this paper.