

What's New About Fake News?¹

Jessica Pepp, Eliot Michaelson, and Rachel Katharine Sterken

Penultimate draft of an article to appear in the Journal of Ethics and Social Philosophy. Please cite the published version.

Abstract

The term "fake news" ascended rapidly to prominence in 2016 and has since become a fixture in academic and public discussions, as well as in political mud-slinging. In the flurry of discussion, the term has been applied so broadly as to threaten to render it meaningless. However, in an effort to rescue our ability to discuss—and combat—the underlying phenomenon that triggered the present use of the term, some philosophers have tried to characterize it more precisely. A common theme in this nascent philosophical discussion is that contemporary fake news is not a fundamentally new kind of thing, but merely the latest iteration of a phenomenon that has played out in different ways across the history of human information-dissemination technologies. While we agree with this, we argue that newer sorts of fake news reveal substantial flaws in extant understandings of this notion. In particular, we argue that no deceptive intentions are necessary for fake news to arise; rather, fake news arises when stories which were not produced via standard journalistic practice are treated as though they had been. Importantly, this revisionary understanding of fake news allows us to see that there is something very different about the new fake news: for stories to be treated as though they were produced by standard journalistic practice, an infrastructure of distribution is necessary. That infrastructure is now fundamentally different than it used to be, raising the possibility that successfully combating the new fake news will require our taking account of this shift.

1. Introduction

On March 4th, 2016 the *National Enquirer* reported that Justice Scalia was murdered by a high-end prostitute employed by the CIA. On October 30th, 2016 a Twitter account purportedly belonging to a Jewish New York lawyer, but in fact run by a white supremacist, claimed that the NYPD had unearthed emails confirming the existence of a pedophilia ring run by prominent Democratic lawmakers—out of a pizza joint in Washington D.C. Thus began “Pizzagate”, a story which spread from Twitter and 4chan to a variety of right-wing websites and, eventually, the Turkish mainstream media.

¹ The authors would like to thank the organizers and participants of the Iceland Meaning Workshop, where an earlier draft of this paper was presented. Special thanks to Elmar Unnsteinsson and two anonymous referees for this journal for valuable feedback. This research was supported by Norwegian Research Council SAMKUL network grant 285131.

² All authors contributed equally to this work.

Both of these are recent paradigm instances of “fake news”. Yet, as we shall see below, there are some significant differences between the two cases. The former is largely contiguous with a type of fake news that has been around for some time. It is produced by organizations resembling real newsrooms, many of which employ real reporters and do occasionally break real news (the *National Enquirer* really did break the story about John Edwards’ extramarital affair, for instance), and is distributed in tabloid form and via websites. The latter is a new sort of fake news, which is interestingly different in a number of ways: (a) it is often produced by individuals; (b) it is often distributed entirely via social networks like Twitter or Facebook; and (c) it relies for its spread not on any sort of physical infrastructure, but rather on the function of those networks, specifically via “sharing”.

In this essay, we will pursue two interconnected goals. First, we will develop a preliminary understanding of what fake news is. On our account, neither deceptive intentions on the part of originators, nor deceived beliefs on the part of consumers are essential to fake news. We will compare our account with recent efforts by some other philosophers (Rini 2017, Gelfert 2018, and Aikin and Talisse 2018), all of whom have argued that characterizing fake news requires some appeal to the intentional state of the agent or agents who produced the content in question. Against such views, we’ll argue that what newer types of fake news (like Pizzagate) reveal is that fake news can arise without any of the efforts or expectations to deceive or mislead that these philosophers have taken to be necessary. In fact, we will argue that attention to the “new” fake news reveals how even older forms of fake news *could* have been produced in this sort of way, even if they actually were not.

Our second goal will be to try to better understand certain features of the contemporary speech landscape that we take to be characteristic of the new fake news. In particular, we think that the speech act of “sharing” is characteristic of, if not essential to, the new fake news. We will explain why we take this to be a specific sort of speech act, how we think it differs from the types of sharing that people were able to partake in before the advent of social media networks, and why all of this matters. In a word, what consideration of this new act of sharing reveals is that the *infrastructure* essential to the spread of fake news (and thus, on our account, essential to fake news itself) has shifted radically in the past few decades. That, in turn, raises the question of what sorts of responsibility ought to be borne by the corporate social networks which have effectively become, structurally speaking at least, the publishers of fake news.

A brief note on what this essay is and is not: importantly, it is not an attempt to capture all the diverse, and very often politically loaded, ways that the term “fake news” is used in contemporary life. Rather, what we hope to do is capture something of the essence of a social phenomenon that has been around for some time, but came to particular prominence around the time of the 2016 U.S. Presidential election. The project is partially descriptive and partially ameliorative; we hope to offer a way of thinking about this phenomenon that both captures a significant amount of what motivated people

to start talking about “fake news” around this time, and which could be useful for practical consideration regarding potential protections and regulations of our media environment. In other words, we hope to better understand a certain fairly coherent social phenomenon that we take to often be picked out by the term “fake news”, one which we think manifests in importantly different ways in different sorts of media environments, and to offer a way of thinking about this phenomenon that will prove useful in efforts to grapple with its political and ethical dimensions and effects.³

2. What is Fake News?

We propose to define fake news as follows:

Fake news is the broad spread of stories treated by those who spread them as having been produced by standard journalistic practices, but that have not in fact been produced by such practices.

This definition locates the “fakeness” of fake news in its status as *news*. Fake news is treated as news by those who spread it, but because it is not produced in the manner in which real news is produced—i.e., by standard journalistic practices—it is not real news. Four important components of this definition need further elaboration: *story*, *broad spread*, *standard journalistic practices*, and *treating as*.

The requisite notion of *story* is quite broad: it encompasses any narrative or report of events, real or imaginary. *Spread* is a matter of how many people read the story and treat it as news. One or two people treating a story as having been produced by standard journalistic processes, when in fact it was not, does not seem to be enough for that story to be fake news. Fake news requires a *broad spread* of a story. The question of how broad is broad enough for fake news is hard to answer precisely. The required breadth would seem to vary with the field of interest for the story. For instance, if a story concerns local politics in a small town, the field of interest for that story—the group of people who are likely to care about the events it reports—might be just a few thousand residents of the town. For this story to become fake news it might only need to be read and treated as news by fifty or so people. By contrast, for a story of national or international interest to become fake news would require a much larger circulation.

It will not be among our aims here to give an analysis of *standard journalistic practices*. We take journalism to be a type of activity focused on (though not defined by) informing members of a society about events relevant to them as members of that society. Journalism is a dynamic social institution with gradual historical origins (mainly in Europe after the

³ We do not mean to commit ourselves here to the claim that the use of the term “fake news” is essential to carrying out this project. As Habgood-Coote (2018) has argued, it might well be that, to combat the sort of phenomenon we are interested in, we would do best to stop calling it “fake news”. While we ultimately disagree with this assessment, we will have to argue against it elsewhere.

invention of the printing press) and forerunners. (See McQuail 2013, chapter 1 for a brief international history.) When exactly this social institution came into existence and where the boundaries of the concept of journalism lie are difficult questions that we will not try to answer. Nonetheless, certain practices do seem to be at least partly constitutive of the activity of journalism; it is these that we refer to by “standard journalistic practices”. We will say more about them in the next section.

Finally, the notion of *treating* a story *as* having been produced in a certain way plays an important role in our understanding of fake news. In general, treating something as having a certain feature is not the same as believing that it has that feature. Often, we treat things as having certain features for certain purposes, even though we do not really believe that they have those features. So, for instance, one might treat a hobby as a business for tax purposes, because the activity seems to meet enough of the IRS’s criteria, while believing that it is still really more of a hobby. Or, one might treat a dear family friend as a family member for the purposes of Christmas gift giving, without believing that the person is actually a family member. Another way of treating something as having a certain feature is to behave as if it has that feature. This also does not require believing that it has that feature. For instance, one might treat a comment as a compliment by responding with a smile and “thank you, that’s nice of you to say,” even if one believes the comment was an insult. Still another way of treating something as having a certain feature is to decide to make that feature prominent in one’s mental categorization of the thing. For instance, the prolific author Harlan Coben said in a 2015 interview, “Every successful author I’ve ever known still has to treat it as a job.” But in the very same interview he also said, “I don’t want a real job; that’s why I became a writer.”⁴ What Coben seems to convey with these two claims is that if you want to be a successful writer, you have to categorize writing as a job in your mind, so that you will do it regularly and diligently. This is consistent with failing to believe that it is, in fact, a job.

There is clearly some overlap in these ways of “treating something as having a certain feature”, and they often go together. (E.g., in acting as if a sarcastic comment is a compliment, one may also mentally categorize it as a compliment, and one may treat it as a compliment for purposes of the conversation.) Treating a story as having been produced by standard journalistic processes, i.e., as genuine news, could be done in any of these overlapping ways. One might feel justified in citing the story during political arguments, thereby treating it as genuine news for purposes of political argumentation. One might include it in a dinner table description of the day’s world events, thereby acting as if it is genuine news. One might focus on its status as news when using it as motivation to get one to take various kinds of political action, thereby making that feature prominent in one’s mental categorization of the story. None of these requires believing that the story is genuine news, i.e. that it was produced by standard journalistic practices.

⁴ <https://www.theguardian.com/books/2015/mar/24/harlan-coben-every-successful-author-still-has-to-treat-it-as-a-job>

This is an important feature of our definition, because it is not clear to what extent those who consume and spread fake news stories actually believe them, or actually believe that they are instances of genuine journalism. It may be that the impetus to treat fake news stories as having been produced in the standard journalistic way comes more from their ability to serve as weapons against perceived political or social enemies, or as solidifiers among political allies, than from readers' belief in their status as genuine news.⁵ Further, it may be that those who consume and spread fake news have an erroneous or premature understanding of what constitutes news and journalism in our new technological and communicative reality. In this case, consumers and spreaders of fake news may believe that certain stories count as news or as produced by genuine journalism, even though what counts as news and genuine journalism in the age of social media is yet to be determined or negotiated by society at large.⁶

Another noteworthy feature of our definition is that it does not require that fake news stories be substantially false.⁷ As mentioned above, the definition locates the fakeness of fake news in its status as news—understood here as the status of having been produced by standard journalistic processes—rather than in its content. Fake news is not genuine news because it has not been produced in the way that genuine news must be produced. Typically it is also false, but this need not always be the case. Just as real news can turn out to contain errors—due either to honest mistakes or to minor breaks in journalistic standards—so too can fake news occasionally (and most likely accidentally) turn out to be true.

Finally, our definition does not require any intentions to deceive or mislead on the part of those who originate or spread fake news. This makes it quite different from several other definitions recently advanced by philosophers. We will take a quick look at three of these. All are attempts to characterize “fake news” in such a way as to capture both the more contemporary phenomenon often referred to by that name as well as most of the work by publications like *The National Enquirer* or *The Weekly World News*. We will continue to treat the Scalia example as a paradigm case of this latter form of fake news and Pizzagate as a paradigm case of the new fake news. Some other instances of the “old” sort of fake news would include: anything from the *Weekly World News*' Bat Boy series or the *National Enquirer*'s 2015 story alleging that Hillary Clinton's deleted emails contained evidence that she is a lesbian. Some other instances of the “new” sort of fake news would

⁵ This is related to the phenomenon of “blue lies,” or lies that are purportedly told to benefit one's group. (See Fu et al. 2008.) It has been suggested, for example, that supporters of Donald Trump do not see his frequent lying as bad because they take it to be done to promote their political side against enemies. (See, for instance, Jeremy Adam Smith's 2017 *Scientific American* blog, cited in Wikforss 2017. <https://blogs.scientificamerican.com/guest-blog/how-the-science-of-blue-lies-may-explain-trumps-support/>.) In a similar way, those who spread fake news stories may view it as good to treat the stories as genuine news, even if they don't believe that they are, because doing so is favorable to some “side” they are on.

⁶ We return to this issue in section 5.

⁷ This is in contrast to (e.g.) Rini 2017, Alcott and Gentzkow 2017, and Gelfert 2018. See also Collins English Dictionary 2017.

include: Trump's tweets about the size of the crowd at his inauguration⁸ or the 2016 claim by YourNewsWire that 25 million fraudulent votes had been cast in favor of Hillary Clinton.

The first definition of fake news that we will consider comes from Rini (2017):

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 known by its creators to be significantly false, and is transmitted with the two goals of being widely re-transmitted and of deceiving at least some of its audience. (E-45)

So, the key elements of Rini's definition of fake news stories are (a) that they purport to describe real world events, (b) that their content is known by their creators to be false, and (c) that they are created with the goal of being widely disseminated and deceptive to at least some readers.

Here is another definition, from Gelfert (2018):

Fake news is the deliberate presentation of (typically) false or misleading claims as news, where the claims are misleading by design. (108)

Gelfert further explains that for a claim to be "misleading by design" is for it (a) to succeed in actually misleading a relevant audience, and (b) *either* to be intended to be misleading in virtue of its specific content *or* to have been produced deliberately using "a process of news production and presentation that is designed to result in false or misleading claims [being spread]" (111).

Finally, a third definition comes from Aikin and Talisse (2018):⁹

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

⁸ This is a particularly interesting example in light of the comment in the text above to the effect that standard journalistic processes, and what counts as journalism, may be in flux due to changes in media. Trump's tweets in this case might be seen simply as assertions by an individual speaker, with no connection, pretended or otherwise, to news or journalistic processes. However, these tweets have been treated as paradigm instances of fake news. This is probably because tweeting by government officials (U.S. Presidents, in particular) can have a semi-journalistic status: just as White House press briefings deliver news to the public (via the press), so White House tweets can deliver news without the press as middleman. There seems to be a similar status for corporate press releases. The issuing of these documents is something the press can report, but the documents themselves are already written as news—they are purpose-built to be *treated as having been produced by standard journalistic practices*. But whereas old-style press releases were also designed to be reported as events—the making of assertions by a government or corporation—in keeping with standard journalistic practices, today's tweets go direct to the public. This might make it easy for such stories—including Trump's tweets about the inauguration—to become fake news.

⁹ Similar definitions can also be found in Levy (2017) and Lazer et al. (2018). See also Mukerji (2018) for a nearby variant which characterizes the relevant intentions not as intentions to mislead or deceive the consumer about the content of the relevant stories, but rather to deceive her regarding whether those stories constitute bullshit.

exploiting their vulnerabilities (cognitive and otherwise), all with a view towards facilitating some decidedly political objective.

All three definitions make the intentions of the originators of fake news essential to the phenomenon. Rini requires that the originators have the goals of wide dissemination and deception of at least some readers. Gelfert requires that the originators intend to be misleading, or at least that they deliberately use processes that are designed to produce misleading claims.¹⁰ Aikin and Talisse require that originating institutions are designed (and so, presumably, intended by someone or some institution) to identify and promote antecedent audience biases. To be sure, many actual instances of fake news have involved such goals and intentions. But it seems to us that neither intentional deception, deliberate use of a process designed to result in false or misleading claims, nor the designing of institutions to feed public biases should be built into the definition of fake news.

Here are some reasons why. First, consider once more the Pizzagate example. For all we know, the white supremacist who instigated this news story actually believed that Clinton's emails contained an elaborate code for talking about pedophilia. "Cheese pizza," for example, was claimed to stand for "child pornography" in virtue of the phrases sharing common first letters. Absurd as this may sound, it is certainly possible that someone could genuinely believe it. Suppose that the progenitor of this story did, in fact, believe in this elaborate code. What's more, suppose that he believed that he had broken the code, and could thus decipher the real contents of Clinton's emails. Working alone as he was, there was little to stop him from believing such things and putting the purported fruits of his labor online. All of this could easily have been completely in earnest, and without any intention to deceive, mislead, or exploit biased thinking. On the contrary, our hypothetical author would likely have taken himself to be *getting the truth out*. We are not inclined to think that, once the story gets up and running, it will fail to count as fake news simply because it was conceived of in a completely sincere manner. Rather, one of the real worries about the media environment we now find ourselves in is that it makes such well-meaning, but highly misinformed, acts of authorship much more likely to gain a significant and wide readership.¹¹

¹⁰It is worth noting that Gelfert's explanation for why he opts for the claim that fake news must be misleading by design does anticipate some of our own thinking. In particular, as Gelfert stresses, he is concerned with the ways in which contemporary fake news is often designed to take advantage of how people will react, systemically, to items shared via social media, e.g. by playing to their cognitive biases (111-13). If we understand him correctly, this means that it would be open to Gelfert to claim that his definition is satisfied when someone aims to take advantage of cognitive biases to spread beliefs that happen to be false or misleading, even if the person doing the spreading actually believes those things to be true. Since Gelfert never explicitly considers such a case, we are not entirely certain whether he would be happy with this outcome. Regardless, we still disagree on two important points: first, about whether fake news needs to actually be false, or even misleading; and, second, about whether it requires any sort of intention to exploit a systemic feature of e.g. social media to spread the claim.

¹¹ Mukerji (2018, section 5) briefly considers a case like this one, which he calls "PIZZAGATE*". Recognizing that this case constitutes a potential counterexample to his theory, Mukerji contends that one can in fact bullshit unwittingly. But this alone is not enough, for on Mukerji's analysis producers of fake

Second, having seen how individuals are capable of generating pieces of fake news without any intentions to deceive or mislead, we are in a better position to see how whole organizations can do much the same sort of thing. Consider an organization like *The National Enquirer*. Suppose that the motivation of the proprietors of this organization is to maximize profit, and that 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. It might just happen that putting forward patently absurd stories turns out, after some trial and error, to be the winning strategy for pursuing (a) and (b). So each person does their part, without anyone intending any deception. In fact, we might even stipulate that everyone in the organization takes the product they are generating to be so clearly unbelievable that they genuinely believe that no one will take the stories appearing in it at all seriously. Still, we submit, this is fake news.¹²

We do not see an easy way of modifying any of these earlier definitions to account for either the recent or the more traditional phenomenon. That said, we think that such a repair strategy—which might involve requiring that originators of fake news either intend to produce misleading stories or *ought to know that this will result from their actions*—is more plausible for traditional sources of fake news like *The National Enquirer* than for recent sources like lone tweeters. This is because these traditional sources have a more robust agential structure—with a CEO, publisher, board, editors, etc.—which might both engage in reflective decision making and be subject to various norms regarding testimony and the public space. On Twitter and Facebook, some of the most prolific generators of fake news have been individuals generating fake news for one country from another, often not in their native language and with little to no reflective oversight. In such circumstances, it is far more plausible that the only real motivations are monetary remuneration, and that these agents’ intentions and predictive awareness of results track those thin motivations. This isn’t even to countenance the possibility of bot-generated fake news—which could very plausibly end up being generated by a machine learner designed only to maximize views in order to promote a certain product, with nothing about deception or mimicry of journalistic practices anywhere in its initial programming,

news must also intend to cover up the fact that they are bullshitting. Presumably, then, Mukerji must accept that we can unwittingly intend to cover up the fact that we are bullshitting. While we cannot offer a full-fledged argument against this claim here, we would note that we find the possibility of such complex unwitting intentions rather far-fetched. We take it to be a point in favor of our theory that we do not have to posit any such unwitting intentions in order to account for a core instance of contemporary fake news.

¹² This kind of concern is also raised in Talisse (2018), though it would seem to cut against Aikin and Talisse (2018)’s requirement that producers of fake news work “with a view towards facilitating some decidedly political objective.”

nor even ever contemplated by its programmers as a possible strategy for view maximization.¹³

3. Standard journalistic practices

Fake news is not just widespread sharing of any old stories, but of stories treated as products of standard journalistic practices that are not in fact such products. As we said above, we are not going to give necessary and sufficient conditions for a story to be the product of standard journalistic practices. But we take it that the relevant standard derives from historical and evolving ideals for the practice of journalism, such as those laid out by Bill Kovach and Tom Rosenstiel in the now-classic book *Elements of Journalism*. These include a commitment to the interest of citizens (or other members of a society); careful verification of putative information received; historical and social contextualization and synthesis of information;¹⁴ active seeking out of information that is or will be important for members of society; transparency about sources, conflicts of interest and unknowns; independence from those covered and from sources; and proportionality in the amount of coverage given to events relative to their likely significance. To qualify as genuine journalism, a process of story production need not maximally fulfill all such ideals. When news production processes live up to these ideals to a significant extent, their outputs have a strong claim to being genuine journalism. But when such processes are substantially out of line with these ideals, they have a much weaker claim to being genuine journalism. This is the case for the usual examples of fake news that we have mentioned. These are stories for which putative information has been fabricated rather than gathered, where there have been no efforts to verify the putative information or to put it into context, and where there is little or no transparency regarding the sources of the putative information or who is responsible for verifying it.

It is important to note that following the standard journalistic practices just described offers no guarantee of reliability, and news sources following standard journalistic practices may well be untrustworthy. For instance, a news organization may, in spite of its best efforts, be fed a line by government sources with their own agendas and little regard for the truth.¹⁵ Should there be little additional information to go on in order to

¹³ It might well be that, just as sentencing algorithms rapidly made use of racial data without their designers ever contemplating that they might do so, view-maximizing algorithms might search the available space of options and conclude that news-like content is the way to go.

¹⁴ One aspect of synthesis is the idea that journalists and news organizations should aim to synthesize information from a politically neutral perspective. To what extent this is actually achievable is an interesting question. However, even if it isn't, there are questions about how much effort towards neutrality is needed in order to satisfy journalistic standards and for journalistic output to count as genuine news. How much lack of neutrality is needed for a story to count as "slanted" or "spun"? How much slant or spin can a journalist or news organization get away with before their story should be counted as fake news? It is worth noting that our account provides the structure to answer such questions about the relationship between spun or slanted news and fake news.

¹⁵ Arguably, this is what happened with Judith Miller's reporting on Iraq for the *New York Times*, circa 2002.

fact-check this line, even well-intentioned reporting can become highly untrustworthy—particularly where reporters are relying on sources who have proven accurate in the past. We take it that such news is hardly “fake” in the sense we are after. Rather, it is the result of the epistemic limitations of standard journalistic practice, which cannot guarantee anything like perfect reliability.

Equally important to stress is that these standards are background expectations for journalistic writing, not ways in which audiences in general would articulate what the standards for journalistic practice are. People identify journalistic writing—that which is subject to these standards—by superficial features such as the use of headlines (with their special grammatical rules), bylines, and certain prose styles and diction. (We will not try to enumerate the latter.) Journalistic writing is also associated with the broader kinds of looks that are customary for print, online, or video news sources—including, increasingly, Twitter feeds.¹⁶ Hence, by employing these looks and superficial features, content originators may produce stories that look like genuine journalism, whether or not they are, and whether or not content originators are aware that they are making their stories look like genuine journalism.

In addition to clear cases where stories are produced in ways that qualify them as genuine journalism, and clear cases where stories are produced in ways that do not qualify them as genuine journalism, there are likely to be plenty of unclear cases. Especially if work in a news organization is divided up, it is conceivable that one journalist might make up a story which others in the news organization sincerely attempt to verify, synthesize, and distribute. Likewise, information might be legitimately gathered by a news organization, but distributed directly without the appropriate process of verification and synthesis. It is also possible that information might be properly gathered, verified, and synthesized, but distributed in a way that obscures or removes responsibility for the journalistic process and transparency of sources. It is not clear that any one of these breaches, combined with a broad spread of the story, is sufficient to make the story fake news. If a rogue reporter at an established news organization fabricates an event which is then verified and synthesized in good faith by others in the organization, it is perhaps not obvious that its spread on social media creates fake news. The same uncertainty applies if news organizations hastily distribute unconfirmed reports during a crisis without appropriate caveats. Finally, an unknown but principled journalist who sets up his own website to look like an established newspaper in order to get people to read his (excellently sourced, well researched) stories may not generate fake news even if his stories get widely spread.¹⁷

¹⁶ On Twitter, this “look” might amount to the mere inclusion of certain hashtags or handles, typically in conjunction with Twitter’s own blue checkmark indicating a “verified” account.

¹⁷ This scenario is to some degree resonant with the real-life I.F. Stone, though Stone was probably working alone out of necessity rather than by choice, and not impersonating real newspapers, but rather producing his own. Still, this raises the question of what to say about journalistic practices which are, intuitively, superior to our present practices—or, in a word, journalistic innovation. Stone’s own work was innovative,

All this suggests that it may not be right to treat the (even partial) fulfillment of any of journalism's characteristic ideals as a necessary condition for a story to escape becoming fake news if it is spread widely. Rather, we take the substantial realization of these ideals to be characteristic features of genuine journalism, any of which might nonetheless be missing in some particular bit of genuine journalism.¹⁸ Of course, if many of the ideals are substantially unrealized on a given occasion, that will be fairly damning and the resulting story is likely to constitute a potential instance of fake news. There are likely to be any number of borderline cases, however, and which side of the border a case sits on may depend not just on the fact that a standard practice has not been followed, but the specific ways in which it has not been followed. Still, in spite of this admitted lack of a clean philosophical analysis, we hope that the reader will have enough of a feel for what we are calling "standard journalistic practice" that she will not object to our invoking it in our attempts to better understand the nature of fake news.

4. What's new about fake news?

We argued above that producers' intentions to deceive, or to make use of production and distribution systems designed to be misleading, are inessential to fake news. Any story or claim that is spread widely and treated as having been produced by standard journalistic practices, although it was not in fact so produced, can be fake news, regardless of its producer's intentions. As we have shown, it is not hard to imagine that prominent examples such as Pizzagate did not involve any intentions to deceive, mislead, or utilize inherently misleading procedures (for all we know, this may even have actually been the case). When social media users are able to easily share articles with large

but shared much of its basic pattern and ideals with ordinary journalistic practice; we suspect that at least some other innovative journalism will follow this same basic pattern, and will thus be recognizable as such. It will be non-standard, but in a way that shares at least a substantial number of features with more standard instances of journalistic practice. We are happy to allow for such superlative instances of journalism to fall under "standard journalistic practice" on our way of talking about things.

¹⁸ As a referee points out, however, one can easily imagine an algorithm capable of accurately predicting—more accurately than current reporting practices, say—events in far-off and inaccessible locales in real time when fed data about surrounding areas. The outputs of such an algorithm might be widely spread and treated as news, and they would not have been produced by standard journalistic practices—for instance, no verification, contextualization or commitment to citizens' interests seems to be involved. But it may seem odd to classify true, reliably generated stories as fake news, as our definition seems to predict. Nonetheless, we think this is the correct verdict about such a case. The matter is complex and deserves more discussion than we can give it here, but our basic reason for the verdict is this: when a story is treated as the product of standard journalistic practices, it is treated as giving a certain kind of epistemic status to the information it contains. This is not merely a certain level of reliability, but also that the reliability itself is transparent to a significant degree. Journalism lets us understand not only how things are but why we should believe that they are that way. Thus, two characteristic features of journalistic epistemology are: (a) a heavy reliance on information obtained non-inferentially, either via testimony, the public record, or photographic or video evidence, with these sources being clearly disclosed; and (b) a lack of reliance on appeals to broad laws, be they of physics, chemistry, or one of the special sciences, to fill in informational gaps when and where such gaps arise. The imagined algorithm does not create the same kind of epistemic relation between readers and the information it conveys. So if its outputs are widely treated as doing so, this is epistemically problematic.

networks of followers who can, in turn, easily share them with their own large networks, even stories produced sincerely but non-journalistically can quickly turn into fake news.

Does the inessentiality of originator intentions distinguish the new, social-media based fake news from traditional fake news? We do not think so. Because traditional fake news tends to be produced by organizations that, in one way or another, do intend to produce at least a certain proportion of misleading stories, it is harder to imagine cases where no such intentions are in place. But in some cases of traditional fake news, individuals, rather than organizations or institutions, have exploited existing journalistic structures to spread claims that they knew to be false or misleading. These cases make it easier to see how even more traditional sorts of fake news could have been, and possibly even were, produced and distributed without any kind of deceptive intentions or misleading design being involved in the process.

An example is Benjamin Franklin's (1782) *Supplement to the Boston Independent Chronicle*. This was a forged "supplement" to a real newspaper, which Franklin printed on his personal printing press in Passy, France. One article in the supplement purported to be a letter from a New England militia captain reporting that a capture of British military goods had yielded two bags of colonists' scalps that the British had obtained from Iroquois and Seneca warriors whom they had put up to the task. According to the letter, these bags of scalps had been carefully catalogued together with grisly descriptions of each colonist's mode of death, and were prepared to be sent back to England. Franklin sent his forged supplement to several friends internationally, hinting slightly at its fakeness and also suggesting that it should make its way into the British press, which it ultimately did. It is debatable whether Franklin's supplement was intended as satire, propaganda, or both. But when the "bag of scalps" story gained traction in the British press of the time, it was clearly fake news. Now, Franklin probably intended to be deceptive and misleading (at least to the average British reader), but it is perhaps not so hard to imagine a different situation in which he created the forged broadsheet for his own entertainment, and in which it fell into the hands of someone who took it for a real newspaper supplement and circulated it to foreign presses.¹⁹ This would still have been fake news. Hence, we maintain that intentions to deceive, mislead, or use production and distribution processes designed to be misleading are not essential to fake news of either the new or traditional varieties.

How then is new, social media-based fake news different from traditional fake news? The difference lies in the relevant media infrastructure, the mechanisms by which fake news is spread. Traditional fake news is spread primarily by institutional news producers, such as the *National Enquirer* or, in Franklin's case, the various British papers. By publishing these stories in their pages, these producers spread the stories to the readership of the publications. New fake news, by contrast, is spread primarily by social

¹⁹ This is not so hard to imagine since Franklin seems to have assembled and used his personal printing press partly as a passionate hobby to provide downtime from diplomacy. See Mulford (2008).

media sharing.²⁰ An originator of new fake news does not bring fake news into existence simply by publishing a story on a web page or twitter account. Rather, social media users collectively bring it into existence by sharing the story with their networks, members of which in turn share it with their networks, and so on.

A similar kind of spreading mechanism is possible for traditional fake news. Franklin shared—in the ordinary, non-social media sense—his forged broadsheet with his friends. If they had also had hobby printing presses, they might have reproduced the forgery and shared it in turn with their friends. And so on. The story might have been spread fairly widely, and treated as news, without any established newspaper or magazine ever publishing it. Perhaps more realistically, a similar forger living in the age of widespread photocopier access might share a forged news article with many friends who could in turn copy it and share it with many more.

Social media sharing drastically increases the ease and scale of such transmission. This explains why the sharing-based spreading mechanism typical of new fake news was more or less unheard of until quite recently. But the difference between sharing something with someone in the general, not specific to social media sense, and sharing something on social media, is not just a matter of ease and scale. For one thing, sharing in the ordinary sense is a broad category comprising many different kinds of actions. A can share a news story with B by telling her about it, showing it to her, or giving her a copy of it. These acts, in turn, can be done in myriad different ways. The specific behaviors people use to convey to each other that a given item is something of interest that they would like the other person to look at and take an interest in are highly varied and idiosyncratic.

By contrast, social media sharing is a comparatively narrow category of action that can be carried out in a limited number of ways. These ways correspond, roughly, to the different operations on social media platforms: one can do a Facebook share, a Twitter share, an Instagram share, etc. Modulo minor platform-specific variations, everyone does

²⁰ Social media has facilitated a change in the mechanisms of news dissemination and uptake which may be thought to mimic those of gossip or rumor. As a result, there may have been a substantial blending of the social functions of the two. This may be, in part, what underlies uneasiness about fake news—that (at least certain sorts of) news and journalistic practice nowadays is more like or serves the function of gossip in many respects. The social functions of gossip are taken by many sociologists to be social bonding and the maintenance of social groups, networks and relations. In addition, gossip is one of the tools that subordinates use to “facilitate open criticism, threats, and attacks” while “shielding their identity” in fear of retaliation (James Scott, 1990). It is also thought to be the mechanism by which we spread information about who can be trusted within a social network, e.g., to explicitly control and punish “free-riders” (Dunbar, 2004). Finally, gossip offers us a way to increase our social capital within a group or network. It is worth noting that these functions are all markedly different from the social and epistemic functions of news, and are plausibly nothing like whatever functions underlie our journalistic practices. Hence, even if *sharing* fake news bears certain superficial similarities to gossiping, and even if certain types of news are now generated largely in order to be shared in ways that will reinforce (largely partisan) social groupings, it remains to be shown that fake news is itself merely a type of gossip. Thanks to Elmar Unnsteinsson for discussion.

social media sharing in the same way. This kind of sharing, we propose, is enough of a uniform action to constitute a distinctive kind of speech act.

The distinctiveness of the act of social media sharing is also observed by Rini (2017), though she conceives of sharing rather differently than we do. According to Rini, sharing is an “ambiguous speech [act] that may or may not be testimony depending on as-yet-unsettled communicative norms.” (E-60, note 10) A second key feature of social media sharing that Rini notes is that: “There is something about social media sharing that seems to deaden people’s normal application of consistency-with-the-world filtering on testimony.” (E-49) She suspects that these “bent aspects of social media testimony play a role in the transmission of fake news” (E-49) and suggests that a way to combat fake news is to unbend such testimony by trying to settle its norms in the direction of accountability. (E-55-58)

While we agree with Rini that sharing is not clearly a form of assertion or testimony, even though it sometimes shares certain features with each of these, we see no ambiguity here. Rather, we take it to be far more plausible that sharing simply isn’t testimony, nor is it a special kind of assertion. Sharing, we would suggest, is a *sui generis* speech act available only within the landscape of social media.

One good reason for thinking that sharing constitutes a new sort of speech act, one that is irreducible to others in the more familiar taxonomy, is that it has significantly different sincerity conditions than any other speech act. In fact, it is not clear that sharing has sincerity conditions at all, as opposed to something weaker like aptness conditions.

Suppose, for a moment, that one were inclined to think of sharing as a new type of asserting or testifying. Then acts of sharing should inherit their sincerity conditions from these acts of which they are a type. But this yields some patently incorrect results. For instance, suppose sincerity in asserting or testifying requires that one believe what one is asserting or testifying (Searle 1969, Stokke 2014). One can appropriately share a news story simply because one likes the picture in it, or because one finds the writing style to be noteworthy, with no commitment at all to the truth of the story itself. We take there to be nothing unsettled about this; it’s just that flexibility can sometimes be mistaken for unsettledness. In contrast, consider “liking” something on Facebook or Instagram: with these, one is acting insincerely if one doesn’t actually have some positive attitude towards the post in question. What’s more, with sharing, there is no easy way of paraphrasing what it is that one is asserting when one shares something. With liking, on the other hand, the act can be fairly straightforwardly paraphrased as “I like this.”

If there is anything in our earlier linguistic repertoire that sharing seems most akin to, it would seem to be pointing, or perhaps pointing and saying “This.” We’ll expand on this similarity below, but first we would note that, if sharing were really a sort of asserting or testifying, this similarity would be highly unexpected. Nor do we think that viewing sharing as a sort of advisory or suggestion adequately captures the flavor of this act. At

least as advisories are standardly understood in speech act theory, they are supposed to express the belief that whatever is being suggested to do—here, to look at what has been shared—would be in the listener’s interest (Bach and Harnish 1979, 48-49). But one can aptly share something even when one believes that looking at it will *not* be in the interest of most of one’s social network. For instance, one can share yet another editorial on the unending chaos in the Trump White House with one’s overwhelmingly liberal social network, even though one takes it that reading this editorial is likely to be a waste of time for pretty much everyone involved. One often shares such stories, we take it, as an act of social positioning rather than to actively suggest to one’s network that they should actually read the story. And we are highly skeptical that one is contravening any established norm of sharing when one does so.²¹

Sharing in social media is an easy, readily-available way to make content available to one’s network while simultaneously indicating that it is you who has made it available to them. While obviously contiguous with earlier modes of sharing, the fact that this action has become one of the basic ways we interact with our social environments sets it apart. Whereas before, one could contrive to do something similar to sharing on a social network by harnessing one’s resources, either on- or off-line, now this is one of the standard ways of interacting with pieces of media. In slightly loaded terms, whereas sharing was once a complex action that could be undertaken with dedication, it is now a basic *affordance* of our interaction with the media landscape.

As mentioned above, we see social media sharing as having much in common with another type of offline speech act, *declarative pointing*. Broadly speaking, pointing is a communicative bodily movement that projects a vector from a body part, indicating a certain direction, location, or object (Kita 2003). Declarative pointing is often distinguished from imperative pointing, the former being aimed at bringing the audience's attention to something, and the latter being aimed at causing the audience to do something.

Brinck (2004) argues that the function of declarative pointing is simply to indicate an object to an addressee, i.e. to direct the attention of the addressee to that object. Further goals that one may have in pointing vary widely and, on Brinck's analysis, are not themselves parts of the function of pointing considered as a type of act. Such further goals might include getting attention for oneself from the addressee, gauging the addressee's evaluative reactions to the objects, coordinating with the addressee on some course of

²¹ Thanks to Elmar Unnsteinsson for pushing us here. We would also note, briefly, that the more one relies on traditional speech act theory to try to understand speech in social media, the more one will run into the problem of having to make sense of machine-generated speech acts. So, for instance, bots now regularly share huge amounts of content on all the major social media networks. Yet, plausibly, these bots do not have anything like the sorts of intentions required for these speech acts to count as felicitous, according to traditional speech act theory. It is open, of course, to the speech act theorist to posit that these are all mere pseudo-speech acts. But given their prevalence on social media, we are loath to think that the liking and sharing undertaken by bots is some fundamentally different kind of act than the liking and sharing undertaken by human beings on these networks.

action involving the object, and so on. The role of pointing itself, however, is just to secure the addressee's attention to the object.

We find Brinck's analysis of pointing persuasive, and we are struck by the similarities with social media sharing. The speech act of sharing might well be seen as a new form of internet-enabled declarative pointing. Ordinary declarative pointing exploits contiguity in time and space of a speaker, an object, and an addressee: the speaker and addressee must be able to see each other and both of them must be able to see the object. Social media sharing exploits the connectedness of a speaker, an object of interest (e.g. a story), and her addressees via the Internet: a speaker and her addressees must be connected to each other by a social media network, and the object of interest must have a web address or be embedded in a shared post. Given the former setup, declarative pointing by the speaker functions to focus the attention of the addressee on the object. Given the latter setup, sharing by the speaker functions to focus the attention of the addressees on the object. Like ordinary declarative pointing, social media sharing may be done with a wide range of aims beyond simply focusing the attention of an addressee on an object. One may aim to get the addressees to believe in what the object says, to share the object further on their own networks, to react with praise or disgust to the object, and so on. But sharing itself is purely indicative.

One might object to this analysis on the grounds that people seem to react to sharing as if it conveys at least some level of endorsement, unless this is explicitly disavowed. As noted above, Rini thinks that the norms around sharing are unsettled. She also thinks that, at present, "our accountability conventions seem to tolerate this instability; we may roll our eyes at 'a retweet is not an endorsement,' but we don't (yet) place most embarrassed retweeters in the same category as outright liars or bullshitters." (E-48) If, in fact, there is a collective eye-roll at those who claim not to have endorsed things they share without further comment, this might suggest that sharing is not purely indicative, like declarative pointing, but also functions to convey some sort of positive endorsement of the thing indicated.²²

This is a fair criticism. One option in reply is to treat sharing as a speech act with the function of indicating an object in an at least slightly-better-than-neutral evaluative tone. However, this proposal might also run into trouble with sincerity conditions. Suppose that someone encounters a news article on which they are perfectly neutral: it strikes them as mediocly researched, mediocly written, and of no particular relevance to their own life, but perhaps of enough interest to some people's lives to merit having been written. This person is also feeling worried about having been silent on Twitter recently, so they decide to share this article, perhaps precisely because it seems so neutral and inoffensive. If sharing conveys at least a mildly positive evaluation of the shared item, which this

²² Alternatively, this might merely suggest that we are apt to suspect that there are further intentions behind many acts of sharing, and that the respondent is being disingenuous about these further intentions—even though all of this is unrelated to any constitutive norms of social media sharing. Still, we take Rini's objection here to be a serious and interesting one, even if the evidence she offers for it is probably not beyond question.

person does not feel, then we should judge this person to be insincere in sharing as they do. But this does not seem right. This person wishes to share something—to achieve joint attention to a bit of online content with her followers—and this is something she feels comfortable sharing. Insincerity does not seem like the right diagnosis, reinforcing the point that sharing does not seem to have sincerity conditions. Given this, it does not in itself convey positive (or negative) evaluations.

Another consideration is that sharing by others is itself a central aim of many who produce online content. When one shares a post, article, video, or the like, one is often giving its creator what they want. This might figure into how one's audience interprets one's share. The addressees may think, “this sharer is rewarding the producer of that content, so she must approve of the content.” But this is not always the situation. Sometimes the wide sharing of some online content is very much *not* what its producer wants. (Think of embarrassing selfies that make their way onto the Internet and go viral.) In such situations, we are not likely to interpret an uncommented share as expressing even mild approval: rather, we are likely to interpret it as expressing disgust, *schadenfreude*, or some other negative emotion. This is because the sharer is doing something against what the producer of the content (presumably) wants. This suggests, once again, that sharing in itself does not have anything like a stable evaluative force.

If this is right, it offers an interesting potential explanation of why the advent of social media has made fake news so prominent. First, there is a natural connection between standard journalistic practices and trust. If followed, these practices should, more often than not, yield information we can use and rely on. The information would come from the world (rather than the imaginations of writers), it would be verified and synthesized into a useful form, and we would know who was responsible if these processes turned out not to have been followed. Not surprisingly, then, the surface appearances that correlate with results of such processes tend to inspire trust. When one encounters a story with this kind of surface appearance, especially if it reports an event whose occurrence would give one some form of emotional satisfaction (for instance, if it says that one's most hated politician has done something terrible), it is easy to let oneself trust the story. Even if one does not trust the story, it may be pleasurable to see the satisfying claim *looking like* a product of standard journalistic processes. And if we are right about social media sharing being a purely indicative speech act, further sharing of the story may not seem especially risky. Thus, the recent epidemic of fake news might be an unsurprising result of the interaction between the newly emergent distributional infrastructure of social media, the correspondingly new speech act of social media sharing, and our shifting attitudes towards journalism as an institution.

Moreover, to the extent that much social media sharing is done with the further aim of expressing evaluative and reactive attitudes toward people and institutions, stories treated as having been produced by standard journalistic practices are especially fruitful objects for joint attention. Here is why. People in the same social networks often have similar evaluative attitudes toward music, art, religion, politics, and so on. As with face-

to-face social interaction, it is pleasing to share one's evaluative feelings with others online, and to get evaluative agreement in return. If two friends are eating a meal together, they can happily gripe about a politician or rhapsodize about a musician in non-specific terms. ("He's so clueless," "Her being in power really scares me," "What a sound she had!", "They don't make musicians like that anymore," etc.) But on social media, the same kind of social bonding seems to require continually introducing new objects of attention. One cannot just express dislike for Hillary Clinton, one has to *point* to something that gives new content to that dislike. A negative story that has the appearance of having been produced by standard journalistic practices is an especially good vehicle for providing such content.

Of course, this explanatory suggestion is only a conjecture. Which psychological factors figure in the spread of fake news is an empirical question that we are not in a position to answer. Still, in the same conjectural spirit, a possible objection to the picture just sketched comes from a recent study suggesting that, at least in the US, it is members of far right groups that share by far the most (of what the researchers categorized as) fake news.²³ Given that these groups also tend to be most suspicious of mainstream media, one might doubt that mimicking the surface appearance of such media would be effective with them. However, we suspect that even those who think journalistic practices are not being followed (at all) by the most prominent institutions of journalism still view them as requirements for *genuine* journalism. The characteristic surface features of journalism may not, by these people's lights, typically indicate genuine journalism, but they are associated with it. News is supposed to be produced in a certain way, and this is how stories produced in that way are supposed to look. A story that looks the right way and says (by a given reader's lights) the right thing might be treated as having been produced in the right way without much additional scrutiny.

5. Sharing and Standard Journalistic Practices

According to our definition of fake news, a story, *S*, is fake news just in case: (i) *S* is broadly spread, (ii) those spreading *S* treat *S* as having been produced by standard journalistic practices, and (iii) *S* was not produced by standard journalistic practices. We have claimed that part of what is distinctive about the new fake news is that it is fulfilled by way of sharing on social media, where sharing is itself understood as a distinctive sort of speech act. It is this profound shift in the infrastructure by which the stories constituting fake news are spread that makes the new fake news importantly different from more traditional fake news. Thus, on our understanding, there really is something distinctively new about the new fake news: it can only come into existence with the advent of social media and, in particular, with the speech act of sharing. Without this new sort of infrastructure, the spread of fake news is constrained by access to the means of widespread copying and dissemination—namely, printing presses or industrial printers,

²³ Narayanan et al. (2018)

delivery trucks, vans, and drivers, and access to vendors or vending machines. With this new sort of sharing infrastructure in place, these constraints fall away and fake news can be generated far more easily by far more people and in more sorts of ways than ever before.²⁴ So the new fake news is contiguous with the old, but also distinct in several important ways.

In this section, we highlight two corollaries of this view. These will serve to clarify the proposal and address some potential objections.

Corollary 1: Our notion of news is *diachronic* and *dynamic*. It is diachronic in the sense that news items are treated as objects that exist over time. It is dynamic in that news items can, in surprising ways, change properties over time. In particular, the status of *S* as an instance of fake news can change over time.

According to our view, at time *t*, *S* is not fake news if at least one of the following conditions obtain:

- C1. *S* is not widely shared
- C2. *S* is widely shared but those sharing it don't treat it as the result of standard journalistic processes.

Suppose, for the moment, that *S* meets one of these conditions at *t* and so *S* is not fake news. On our proposal, at a later time, *t'*, *S* can turn into an instance of fake news either by becoming widely shared or by being treated as the result of standard journalistic processes. For instance, a satirical story from *The Onion*, which has many of the trappings of real news, might be shared widely at time *t*, but not be treated as the result of standard journalistic practice. So it will not count as fake news at *t*. At *t'*, however, it might well become the case that the story, still widespread and still intended by its originators as satire, for whatever reason comes to be treated as the result of standard journalistic practices. On our account, this same story will now count as fake news.²⁵

²⁴ Again, to stress a point we have gestured at before, it is not exactly that this new sharing infrastructure couldn't have been mimicked with certain aspects of the older infrastructure. A copy machine and the post office can be used rather effectively to spread fake stories (see Simon Garner's "98 Years of Mail Fraud," <https://www.theatlantic.com/technology/archive/2018/05/98-years-of-mail-fraud/559661/>). Email makes this sort of thing even easier. But even if it is possible to spread fake news by these means—and we have no doubt that it is—something significant has clearly changed with the advent of social media. This is what we tried to capture above in pointing out that social media infrastructure has made sharing into a kind of basic affordance, as easy as (and very similar to) pointing at something in your visual field to achieve joint attention with an audience.

²⁵ There are, in fact, real world cases of something like this, though many involve the simultaneous spread of a story and treatment of it as non-satirical. So, for instance, Iran's FARS news agency apparently republished, word for word, a 2012 *Onion* article claiming that polling suggested that rural Americans preferred Ahmadinejad to Obama; in fact, the *Washington Post* reports that this sort of thing, minus the outright plagiarism, is actually fairly common, and not just in foreign countries (see: https://www.washingtonpost.com/news/worldviews/wp/2015/06/02/7-times-the-onion-was-lost-in-translation/?noredirect=on&utm_term=.05364a138d71). Assuming this story became widespread in Iran, it would now count as fake news in Iran on our story. This brings out another interesting aspect of the view that we will not have time to focus on in the main text: on our view, it is possible for a story simultaneously to count as fake news (relative to some region) and not to count as fake news (relative to some other region).

Essentially, we propose to think of a piece of news in a similar way to how Kaplan (1990) thinks about words: they are objects that exist over time and different stages of a word (story) can have different properties. As such, on our view, a news item can turn into fake news, not by its intrinsic features changing (e.g., the set of sentences that constitute the story or the propositions expressed by those sentences), but by becoming widely shared or by being treated in a new way by those sharing it.

The aim of our notion is in part to shift attention away from a narrow focus on the informational content of a story and the intentions and acts of the original author. These are not the full story, we have argued, about what makes a piece of news what it is. The history of the news story, including the way people treat it on social media, is partly constitutive of whether, at a given time, a particular news story will count as fake news or real news.

Corollary 2: We just articulated the idea that a consequence of our diachronic and dynamic conception is that what, at t , wasn't fake news, can turn into fake news. Our view also entails the contrapositive of this: what once was fake news can turn into not-fake news.

Here is how this might conceivably happen: at a time t , a story, S , is an instance of fake news (i.e., at t , S is widely shared and treated as a product of standard journalistic practices when in fact it is not). Now at a later time, t' , a journalist, J , looks into S . J is a hard-working, thorough and conscientious journalist working on a normal deadline. She does her due diligence—that is, she does what she can, given restrictions of time and resources, to check on the accuracy of S . When she does that, she keeps encountering sources that unbeknownst to her are inaccurate and supportive of S . The bad sources she encounters in her research all confirm S —this might be because she's in an informational environment where fake news is ubiquitous. Keep in mind that standard journalistic practice isn't a guarantee of infallibility: J can follow standard journalistic practice and get it wrong. That's what happens in the case under consideration: J publishes the story and it is shared further. However—and this is the interesting part— S is no longer an instance of fake news, according to our account, since in order to be an instance of fake news, S cannot be the result of standard journalistic practices. So at time t , S was fake news, then at t' (when S is published by J), S is no longer fake news—even though S 's content with all its inaccuracies is constant across t and t' . Alternatively, one is welcome to imagine the same basic scenario, but where S is actually accurate the entire time. That is equally well a possibility, on our view.

Is this a feature or a bug of our proposal? We see it as an advantage of the view. Our view suggests focusing on the historical development of S and thinking of a news story as having different properties at different times. Our proposal adds to this that the presence and absence of standard journalistic practices create crucial junctures in such histories. As a result, our account allows not-fake news to turn into fake news and vice

This should come as no real surprise given that our account of what it is to be fake news is essentially relationalist in nature—as opposed to its rival, more essentialized, accounts.

versa. When fake news turns into not-fake, but nonetheless false news, as in the first variant of the case described above, we can provide an accurate diagnosis of what has gone wrong: *S* is now an example of how standard journalistic practices can fail, without the journalist or news organization being to blame. Journalists and news organizations are fallible and, in the case described, we have a way of understanding what went wrong (*J*'s sources were bad). We are proposing that it is crucial to mark this defect as importantly different from the defect *S* suffers from at time *t*. In the case where *S* turns out to have been accurate the whole time, we see how standard journalistic practice can manage to grant a story real epistemic *bona fides*, even when the causal origins of that story are epistemically suspect. In both instances, real journalism can, on our account, *salvage* a fake story. When that story is false, we grant that it will not have been worth salvaging. But, when it is true, it seems to us that we should want a story of how something of real epistemic worth can be gained; our account provides that story.

That said, this is a feature of the view that would be fairly easy to avoid if one were so motivated. Here is one strategy: we've not given you necessary and sufficient conditions for being the same news story over time. That is, we haven't answered the question: When does a set of sentences (or propositions) count as the same news story? It might be that the individuation conditions for news stories are such that, in a case like the one described, *J*'s research and reporting results in a new story, *S'*, which is not identical to *S*. *S* is a source of *S'*, and *S* is fake news. But because *S'* is the result of standard journalistic processes (albeit in a situation where these have not yielded accuracy and appropriate contextualization), *S'* is not fake news.

So a diachronic and dynamic conception of a news story will need to be supplemented with an account of the individuation and perdurance conditions for news stories. One needs to know what counts as the same news story over time. Our view highlights this question, but doesn't yet answer it. Indeed, on our view, the question isn't settled prior to the development and uptake of new communicative technologies. In the shift from a primarily print, picture and audio-based conception of news to one driven by social media, what counts as news (and so fake news) has changed. The same is likely true of the question of what counts as the same news story over time. The medium, and in particular the distributional infrastructure that it engenders, is in part constitutive of the nature of news and news stories, and is bound to play a role in determining their means of individuation. These are issues we won't settle here, but we note it as an advantage of our view that it highlights the importance of these questions.

7. Conclusion

We have argued that fake news is best conceived of as something that can arise independently of anyone's having intended to produce it. What's more, we have attempted to sketch a way of defining fake news in terms of the way that a story is treated, rather than in terms of how it is created. We think that this definition is superior to earlier ones, not just because it can handle cases where there was no intention to produce a piece

of fake news—and even cases where there was no intentional agent producing the fake news—but also because it highlights the ways in which what will count as fake news can change over time in significant ways. This can happen when journalistic standards change, when the look of news changes, and when the means of distributing news stories change. Stories that once would not have counted as fake news can become such, and vice-versa.

Our definition as it stands is schematic in certain ways, in particular when it comes to the notion of “standard journalistic practices”. We have suggested that this is, in fact, a virtue of the view. It offers us a decent first-pass understanding of fake news, but also allows for a more refined use of the concept as it becomes clearer what it is that we ought to consider standard journalistic practice in the present context. Many contemporary debates on what, in particular, should count as fake news hinge, we take it, on just these questions. So making this explicit, rather than merely tacit, can hopefully serve as a step towards a more useful refinement of this concept—one that might help to clarify how we might reform our social media space in order to preserve its enormous potential to connect people without descending into a realm of spin and disinformation that can pose a fundamental threat to free society.

We will close by highlighting one aspect of our account that *is not* schematic. Unlike its rivals, our account makes widespread dissemination an essential feature of being fake news. While we expect this to be a controversial aspect of the view, we also take it to be a point worth defending. It is worth defending not only because it allows our definition to capture an important set of cases—those which involve no sort of intentions to deceive or mislead—that other views cannot easily account for, but also because it helps us to understand the import of the radical changes in the infrastructure of fake news that we have witnessed in the last decades. These changes, of course, aren’t changes in the nature of fake news on our account. But they are fundamental changes in the ways that nature can be manifested in our world; they represent a huge shift in how fake news can be generated. Accordingly, we take our account to be helpful to a broader sort of ameliorative project to which we are sympathetic: we should hope for our philosophical theories not just to achieve extensional (or intensional) adequacy with respect to some concept. Rather, at least for a certain class of concepts, we should hope for those theories to aid in our attempts to usefully grapple with the world around us.

We think our proposed definition of fake news meets this criterion better than its rivals can. First, it highlights the important role that regulating *unintentional* fake news is going to have to play in creating a healthier epistemic environment. Second, and relatedly, it points to the importance of infrastructure in these debates. The infrastructure of news used to be controlled by publishers, who could be held to account for the things that were reported in their papers. With the advent of social media, this is no longer the case. In the West at least, social media companies have been largely exempted from the sorts of laws that hold publishers to account for the contents of the stories they publish. We would hardly recommend that these networks be subjected, wholesale, to libel or

truth-in-advertising laws of the sort that more traditional publishers were subject to. That said, if constraining fake news is the sort of worthy policy goal that we take it to be, then the possibility of subjecting these networks to greater scrutiny with regard to stories that (a) become widespread, and (b) exhibit the superficial features of real news might be worth considering. For, if we are right about fake news, then such a policy might allow for a significant reduction in the spread of fake news with only minimal impacts on the sorts of individual free speech rights that most of us will—rightly, we think—take to be well worth protecting.

References

- Aikin, Scott and Talisse, Robert. 2018. On “Fake News.” 3 *Quarks Daily*. <https://www.3quarksdaily.com/3quarksdaily/2018/05/on-fake-news.html>
- Allcott, Hunt and Gentzkow, Matthew. 2017. Social Media and Fake News in the 2016 Election. *Journal of Economic Perspectives*, Volume 31, Number 2: 211–236.
- Bach, Kent and Harnish, Robert. 1979. *Linguistic Communication and Speech Acts*. Cambridge, MA: MIT Press.
- Brinck, Ingar. 2004. The pragmatics of imperative and declarative pointing. *Cognitive Science Quarterly*, 3(4), 429-446.
- Collins English Dictionary (2017) Definition of ‘Fake News’, available at <https://www.collinsdictionary.com/dictionary/english/fake-news>.
- DeFleur, M.L. and S. Ball-Rokeach. 1989. *Theories of Mass Communication*, 5th edn. New York: Longman.
- Dunbar, R.I.M. 2004. Gossip in evolutionary perspective. *Review of General Psychology*, Volume 8, Number 2: 100-110.
- Fu, Genyue; Evans, Angela; Wang, Lingfeng; and Lee, Kang. 2008. Lying in the name of the collective good: a developmental study. *Developmental Science Jul; 11(4): 495–503*.
- Gelfert, Axel. 2018. Fake News: a definition. *Informal Logic*, Vol. 38, No.1 (2018), pp. 84–117.
- Habgood-Coote, Joshua. 2018. Stop Talking About Fake News!, *Inquiry*, online first.
- Kaplan, David. 1990. Words. *Proceedings of the Aristotelian Society, Supplementary Volumes*, Vol. 64: 93-119.
- Karlsson, Michael. 2011. The immediacy of online news, the visibility of journalistic processes and a restructuring of journalistic authority. *Journalism* 12(3): 279–295.
- Kita, Sotaro. 2003. Pointing: A foundational building block of human communication. In S. Kita (Ed.), *Pointing: Where language, culture, and cognition meet* (pp. 1-8). Mahwah, NJ: Erlbaum.

- Kovach, Bill and Tom Rosenstiel. 2007. *The Elements of Journalism: What Newspeople Should Know and the Public Should Expect, Completely Updated and Revised*. New York: Three Rivers Press.
- Lazer, David; Matthew A. Baum, Yochai Benkler, Adam J. Berinsky, Kelly M. Greenhill, Filippo Menczer, Miriam J. Metzger, Brendan Nyhan, Gordon Pennycook, David Rothschild, Michael Schudson, Steven A. Sloman, Cass R. Sunstein, Emily A. Thorson, Duncan J. Watts, Jonathan L. Zittrain. (2018) The Science of Fake News. *Science* 09 mar 2018: 1094-1096.
- Levy, Neil. 2017. The Bad News About Fake News. *Social Epistemology Review and Reply Collective* 6, no. 8: 20-36.
- McQuail, Denis. 2013. *Journalism and society*. Sage.
- Mukerji, Nikil. 2018. What is Fake News? *Ergo*, Volume 5, Number 35.
- Mulford, Carla. 2008. Benjamin Franklin's Savage Eloquence: Hoaxes from the Press at Passy, 1782. *Proceedings of the American Philosophical Society*, Vol. 152, No. 4 (Dec., 2008): 490-530
- Narayanan, Vidya, Vlad Barash, John Kelly, Bence Kollanyi, Lisa-Maria Neudert, and Philip N. Howard. 2018. Polarization, Partisanship and Junk News Consumption over Social Media in the US. Data Memo 2018.1. Oxford, UK: Project on Computational Propaganda. comprop.oii.ox.ac.uk
- Rini, Regina. 2017. Fake News and Partisan Epistemology. *Kennedy Institute of Ethics Journal*, Volume 27, Number 2 Supplement, June 2017: E-43-E-64.
- Scott, James C. 1990. *Domination and the Arts of Resistance*. Yale University Press.
- Searle, John. 1969. *Speech Acts*. Cambridge University Press.
- Smith, Jeremy Adam. 2017. How the Science of “Blue Lies” May Explain Trump’s Support. *Scientific American*, Guest Blog. <https://blogs.scientificamerican.com/guest-blog/how-the-science-of-blue-lies-may-explain-trumps-support/>.
- Stokke, Andreas. 2014. Insincerity. *Nous*, Vol. 48, No. 3: 496-520.
- Talisse, Robert. 2018. There’s no such thing as fake news, and that’s bad news. *3AM Magazine*. <https://www.3ammagazine.com/3am/theres-no-such-thing-as-fake-news-and-thats-bad-news/>
- Wikforss, Åsa. 2017. *Alternativa Fakta: om kunskapen och dess fiender*. Fri Tanke.