

Fake News Epistemology

(draft)

Michel Croce and Tommaso Piazza

Introduction

Social media like Facebook, Twitter and Instagram are increasingly shaping our lives. Not only are our social identities—the way we look at ourselves, and how we are perceived by others—more and more mediated by these technologies, but also our intellectual selves—the way we acquire and manage our beliefs—are becoming increasingly dependent on the online environments we inhabit. This consequence has appeared especially problematic in view of the recent upsurge of the phenomenon of *fake news*. Social media and online communities have in fact proven especially hospitable to the proliferation, and wide-spread acceptance, of false or otherwise misleading claims pertaining to the public arena, like alleged reports of political candidates' misconduct, inaccurate figures about the economic impact of political measures, ungrounded predictions about the consequences of certain large scale medical treatments etc. To the extent to which we routinely and increasingly stock up on information online, we are then more and more prone to take up inaccurate beliefs. Hence, the chances that we act in the public sphere—by debating, campaigning, or voting—upon a poor or insufficient factual basis dramatically increase.

This problem is first and foremost located at the individual level, where systematic reliance on online sources shared through social media presents each of us with the threat of *epistemic deprivation*, namely with a significant decrease in epistemically valuable goods (true or justified beliefs) and an increase in epistemically degenerated goods (false or unjustified beliefs). The problem has a clearly social dimension to it, though, as our society at large runs the risk of deliberating about itself and about its own future in a state of systematic disinformation, and hence in a way which is potentially disconnected from the facts. For all these reasons, the phenomenon of fake news has attracted the attention of many scholars from a variety of different disciplines, and is now almost unanimously regarded as a disease for which we (individually, and *qua* members of a democratic society) urgently need a good therapy (Habgood-Coote *forthcoming* is perhaps an exception, as he suggests that the therapy we need is to be freed from the very notion of fake news).

This paper is primarily devoted to review and critically engage with the recent *philosophical* literature about fake news. This literature has primarily concentrated on the nature of the phenomenon under discussion, so on the *definitional question* of what fake news in fact *is*, and on what distinguishes it from related phenomena like lies, rumours, hoaxes, and satire. A second important question, addressed within this literature, pertains to the *mechanisms* that prompt the proliferation of fake news in the online environments that are of special interest to this special issue, as well as in other more traditional off-line contexts. A third final question addresses the normative dimension of fake news; in particular, among others, the question about who is to be

blamed, and from which point of view, in the process originating in the release of fake news, and culminating in their propagation and/or subsequent consumption. The paper addresses these three questions in turn and is organized as follows.

§1 inquires into the nature of fake news. As it will emerge from our review of the extant literature, much of the debate pertains to the properties an assertion should possess in order to be an assertion whose (literal or implied) content is fake news. After individuating four kinds of relevant properties and highlighting how specific accounts of the nature of fake news combine these properties, we will propose our own account of what fake news is.

§2 inquires into the ways in which fake news typically propagates. We start by identifying three key roles: the originator/assertor of a piece of fake news, the propagator of fake news, who receives and shares it over social media, and the final consumer of fake news (the second and third role may of course overlap). We then introduce the most likely causal factors—that we classify as *individual* or *social*, depending on whether they essentially involve reference to one's membership within a given community—that explain the behaviours of those who play the above roles.

Finally, §3 addresses questions of responsibility, and investigates into the normative (mainly epistemic) (de)merits of, especially, propagators and final consumers of fake news. After distinguishing three main answers—ranging from those who consider consumers epistemically virtuous and blame the environments within which fake news proliferates to those who place blame on single individuals and their epistemic vices—we will defend our own approach, which combines the most salient merits of the other two.

§1 – *What is Fake News*

The scholars who have addressed the definitional question above tend to agree about the genre to which fake news should be taken to belong as a species. For Rini (2017: E-44), fake news is a certain type of *information*; for Mukerji (*forthcoming*: 14) fake news is *something* that is asserted; for Jaster & Lanius (*forthcoming*: 2) it is a specific kind of *news*; finally, Gelfert (2018: 103-4) explicitly states that fake news is a *species of information*, a *form of news*, something *that is presented as true*. As is clear, all these characterizations point in the same direction: fake news is the propositional content of certain specific linguistic acts. We'll address in a minute the question about *which linguistic act* fake news is the propositional content of. For the time being, however, we can ask *what* kind of content of a linguistic act fake news is supposed to be. All participants to the debate agree that fake news is often the *literal* content of linguistic acts. One of the best examples of the fake news spread during the 2016 US presidential elections—the so-called pizzagate—was *that* Hillary Clinton was involved in a child trafficking ring, and that she was so involved was literally stated by many posts and fake news reports issued across the internet.

However, it is tempting to conjecture that fake news need not be the literal content of an assertion, and that it is sometimes sufficient that it be the content *implicated* by an assertion whose literal content is not *fake news*. Jaster & Lanius give the following example to illustrate the point. Suppose an online newspaper reports the following story: “After the refugees arrived, 47

burglaries occurred in the village”. Suppose that what this report literally states is true, as exactly 47 burglaries occurred in the village in the period referred to within the report. The report may nonetheless *convey* a piece of fake news, for it certainly implicates that the 47 burglaries occurred *because* of the arrival of the refugees, and that their arrival has caused a raise in the crime rate. And as is obvious, it clearly constitutes a possibility that the 47 burglaries occurred after the refugees’ arrival and that they did not occur because of that, and that the refugees didn’t raise the crime rate¹. Were any of the latter two scenarios to be realized, it would be natural to classify the initial report as fake news, despite the fact that what it literally states is true².

We can now ask what properties turn an assertion into an assertion whose literal or implied content is fake news. As anticipated, the debate has concentrated on four kinds of properties: epistemic properties, intentional properties, sociological properties, and format properties. Let’s inspect them in turn.

The label “epistemic” in “epistemic properties” is meant in a broad sense, as encompassing truth and falsity. Truth and falsity, *stricto sensu*, are semantic and not epistemic properties. However, in the present context they can be legitimately counted as epistemic because they are clearly *related* to our epistemic practices: truth is what these practices aim at, and falsity is what they aspire to stay clear of. *Falsity* is a natural culprit when, in this broader sense, it comes to the epistemic properties of the assertions whose content is fake news. As recorded by Gelfert (2018: 99), Facebook has gone as far as to re-label fake news as “false news”, in order to convey in a more immediate way the key thought that fake news, for them, are nothing but *false* news disguised as genuine ones.

The idea that fake news is often false news is commonly stressed in the literature. According to Rini (2017: E-45), when the creator of a piece of fake news asserts it, they know that its content is “significantly false”; thus, for Rini, fake news must be significantly false, otherwise its creator could not know so. Other commentators adopt a more nuanced stance. Allcott & Gentzkow (2017), Gelfert (2018) and Jaster & Lanius (2018) agree that false belief on the part of its final consumer is the end product of the successful dissemination of fake news. However, they recognize that falsely asserting that *p* is just one possible way to achieve the goal of instilling in one’s addressee the false belief *p*. One may achieve the same goal by making a *misleading* assertion, namely an assertion that may engender the formation of false belief not because it states something false but because, although it states something true, it merely implicates something false. To summarize the point, Jaster & Lanius say that fake news is characterized by *lack of truth*, where this can “come about through the falsity of a news report or ... through its misleadingness” (2018: ???).

Lack of truth in Jaster & Lanius’s broader sense—as false or misleading assertion—is commonly taken in the literature not to be a *sufficient* condition for fake news. Two kinds of cases

¹ Suppose, for instance, that 47 is exactly the number of burglaries one would have expected even before the arrival of the refugees, or that one would have expected a greater number, and the arrival of the refugees has contributed to lower it.

² Mukerji (*forthcoming*: 27) seems to share this intuition to a lower degree, and proposes to classify cases like these, in which a literally true report implies something false, as cases of bullshit journalism rather than cases of fake news.

support this idea. First, honest journalist mistakes may state literally false contents, yet they are typically not equated with fake news. Second, in the attempt to simplify some difficult topic and make it more accessible to a broader audience, a journalist might end up publishing a story which, although it conveys a picture which is broadly correct, is replete with minor falsities (Gelfert 2018: 99). Again, calling the journalist's report "fake news" would run against the common use of the label.

Some controversy, however, surrounds the claim that lack of truth is, as many contend, a *necessary* condition of it. In order to properly assess alleged counterexamples, namely cases of fake news conveyed by assertions that do not lack truth in Jaster & Lanius's sense, it is necessary to discuss the intentional properties that an assertion must possess in order to convey fake news. Hence, we shall postpone a discussion of whether lack of truth or some other epistemic property is a necessary condition of fake news until the end of this section.

By *intentional properties* of an assertion we mean a series of psychological properties that may accompany and motivate the act of asserting a given propositional content. Of particular relevance, in the present context, are the *intentions* that motivate the act of making a given assertion, and the *attitude* with which one makes it. For instance, many contend that in order to assert a piece of fake news one has to have the *intention* of deceiving their audience (Dentith 2017: 66, Gelfert 2018: 108). This suggestion explains why honest journalist mistakes do not amount to fake news: for although they may incorrectly picture some state of affair, they do so *inadvertently*. The suggestion also explains the difference between fake news and certain forms of satire or news parody—like the *Onion* (or, in Italy, *Il Lercio*)—that publish plainly false stories without the intention to deceive, but rather with the goal of mocking real people by mimicking the style and look of mainstream media.

Jaster & Lanius (2018) agree that in many cases the intention to deceive turns a misleading or utterly false assertion into the assertion of fake news (forthcoming: ???). However, they do not think fake news necessarily involves the intention to deceive. On their view, fake news lacks what they call *truthfulness*. Usually, an assertion is said to be untruthful when the speaker believes its content to be false. Jaster & Lanius, however, apparently mean a different thing, namely an assertion that is not made with the purpose of reporting the truth. Clearly, an assertion made with the intention to deceive is not an assertion made with the purpose of reporting the truth. However, an assertion can lack the purpose of reporting the truth—that is, can lack truthfulness in Jaster & Lanius' sense—even if it is not made with the intention to deceive. For this to happen, it is sufficient that the speaker, in asserting a given content, be *indifferent* to its truth. As an example, think of bait farms that produce and disseminate appealing and sensationalist contents for the sole purpose of creating as many clicks as possible (and of gaining money proportionally). The producers of these posts are only concerned with their profit; thus, they are untruthful in the second sense of being unconcerned with the truth of what they report.

To assert a content with no concern with its truth, as Jaster & Lanius explicitly acknowledge, is to utter what Frankfurt (2005) has famously called *bullshit*. Bullshitters, like the

Macedonian villagers above, don't care whether what they say is true or false. They just pick up what they say, or invent it out of thin air, for the reason that it suits their practical ends.

The analysis of fake news as bullshit in Frankfurt's sense is at the centre of Mukerji (*forthcoming*)'s account of fake news. Also Mukerji contends that one has to be indifferent to the truth of one's assertion for its content to be fake news.³ Mukerji further notices that the assertor of fake news, like a Frankfurter bullshitter, must be also willing to deceive its addressees. This does not necessarily involve the intention to deceive about the facts. Rather, fake news assertors may just intend to deceive their addressee about their own attitude towards the truth: that is, they may fool the addressee into believing that they are concerned with the truth of what they are asserting when they may not be. This, as we will see, seems to require that by asserting a given content the speaker falsely implicates that there is some available body of evidence to which they are thereby paying heed. We shall take up again this topic at the end of this section, where we shall connect it to the question about the epistemic properties of fake news assertions.

For the time being, let us address what we have called the sociological properties of fake news. Some theorists contend that a piece of fake news, in order to be so, must actually attain the goal of wide circulation and uptake (Gelfert 2018). Others, like Rini (2018), just require that fake news be transmitted with *the goal* of wide circulation and re-transmission. We tend to favour an irenic solution. We think attaining wide circulation, other things being equal, may be *sufficient* for a message to convey fake news. However, we do not believe that any case in which a message fails to achieve wide circulation should *not* be regarded as a case of fake news. Much, in our opinion, depends on the message's *potential* for wide circulation. In other words, we believe that an assertion can be regarded as the assertion of a piece of fake news—regardless of whether it achieves the goal of circulating widely—provided that it is made over a medium that is at least conducive to the attainment of that goal.

An example by Gelfert can be useful to illustrate this suggestion. He reports that when he was a physics student, a mentally disturbed man used to stand next to the library's photocopier and to make access to it conditional on the acceptance by prospective users of a pamphlet denying Einstein's theory of relativity. Gelfert rightly observes that this pamphlet would not have been (the scientific equivalent of) fake news even if it had possessed the right format, namely if it had managed to mimic in a successful way the "look and feel of a professional journal article" (102, fn. 7). For Gelfert this is due to the fact that the pamphlet failed to realize the goal of widespread circulation (102). However, we believe that Gelfert's example more aptly illustrates a slightly different moral, namely that the way in which an assertion is made must not compromise its chances to circulate widely. Suppose that the pamphlet had been published on a website successfully counterfeiting the look and feel of an online scientific journal and that, nonetheless, it failed to achieve wide circulation. In this case, independently of its lack of success, we would be inclined to characterize it as (the scientific equivalent of) fake news: namely, as something

³ Mukerji contemplates various senses in which one can be indifferent to the truth of what one asserts. A speaker, for instance, can be indifferent to truth in the sense of attaching no value to the truth of her assertion. For Mukerji the assertor of fake news need not be indifferent to the truth of what she asserts in this sense. It is sufficient that she is indifferent to the truth in the weaker sense that she would have made the assertion whether or not it was true.

belonging within the same category as – and much less successful than – e.g. the theory that the Earth is flat.

As noted by many parties, the assertion of fake news must also possess certain format features that make it appear like it is real news. Fake news, in order to be so, must be *presented* as news. This suggestion, to be sure, is rather vague and leaves many questions unanswered. In the present context, we focus on the central question asking *where* a piece of fake news has to be presented as news. Some theorists (e.g. Klein & Wueller 2017: 6) insist that fake news requires online publication. Others, like Rini (2017: E-45), simply take notice of the existence of a strong contingent relationship between fake news and social media. Mukerji (forthcoming: 10) warns against the mistake of misidentifying a contingent relationship for a conceptual connection, and gives some examples of fake news spread across traditional media such as newspapers and the radio.

Our inclination in this dispute is to side with Mukerji (and perhaps Gelfert 2018) and grant that fake news, unlike the label that has been recently coined to refer to it, antedates social media sharing. However, we also think that the new media have played a crucial role in shaping the spread of fake news in its current form, to the extent that it has now acquired peculiar features. On the one hand, as noted by Gelfert (2018: 102), online journalism has made it far easier to get a story published with the hallmark of serious reporting and remove “many of the traditional markers of traditional journalism”. So, if not the phenomenon itself, its *dimensions* have been greatly influenced by the emergence of the new media. On the other, social media have created entirely new business opportunities for bait farmers to explore, by spreading contents designed to be likeable and sharable in the attempt to generate considerable revenues from pay-per-click advertisement. So, above and beyond the dimension of the phenomenon, the online medium through which it is transmitted has profoundly re-shaped the *motives* and *aims* behind it.

Now that we are done with our discussion of the properties of assertions that can be legitimately taken to be assertions of fake news, we can take up again the discussion of epistemic properties, which we left with the question about whether lack of truth—namely utter falsity, or literal truth paired with a false implicature—is a necessary condition for an assertion to be the assertion of fake news. As many theorists recognize (Jaster & Lanius *forthcoming*, Mukerji *forthcoming*), certain possible cases initially invite the suggestion that lack of truth is not a necessary condition. If this suggestion is right—as we believe it is—the defect of the relevant fake news must not be located in its relation to truth, but in some other epistemic shortcoming. In what follows, we shall endeavour to describe which shortcoming it is, and show that the same shortcoming afflicts untrue fake news of a more regular sort.

All cases considered in the literature revolve around *accidentally true* fake news. As an example, suppose *Russia Today* fabricates the story that Hilary Clinton committed tax fraud (Jaster & Lanius *forthcoming*: 10). The story is simply *invented*. The authors of the story wanted to discredit Hillary Clinton as a political candidate and thought her image would have been considerably damaged by the allegation that she didn't pay her taxes. So, in spite of having no reason to believe the story, they published it online. The story goes viral, and millions of American voters become

convinced that Hillary Clinton committed tax fraud. As it happens, this story happens to be true: Hillary Clinton secretly committed tax fraud. Is the publication of the story the assertion of fake news?

Jaster & Lanius maintain that we might feel tempted to look at the *Russia Today* case as involving a piece of fake news, but argue that this temptation should be resisted.⁴ They suggest that the intuitive pull of the verdict that tax fraud *is* fake news vanishes when we carefully distinguish between spreading fake news and merely attempting to do so. In this case, we see that the publication of the story about Hillary Clinton is merely an attempt to publish, and is not the actual publication, of fake news. They argue for this claim by drawing a parallel with lying.⁵ According to them, we do distinguish – and punish accordingly – lying from the attempt to lie. When one attempts to lie, one utters what one believes to be false in the attempt to deceive one’s audience; but when what one asserts is in fact true, and so one did not deceive anyone, one has *merely* attempted to lie, and has not lied at all. In the same way, for Jaster & Lanius when one attempts to spread fake news, one publishes what one believes to be false in the attempt to deceive one’s audience; but when what one has published is in fact true, one has *merely* attempted to spread fake news, and has not managed to do so.

This argument by analogy can easily be resisted. Begin to notice that the claim that lying does not require the statement of a false proposition is widely accepted in the philosophical literature about lying. On the standard definition of lying, to lie is to make a statement believed to be false, with the intention of getting another to accept it as true (Isenberg 1964, Chisholm & Feehan 1977, Primoratz 1984, Williams 2002, Mahon 2008, Lackey 2013). According to this definition, the assertion of a true statement amounts to a lie when the assertor mistakenly believes their statement to be false and makes it with the intention to deceive their addressee. If lying tolerates the assertion of a true proposition, the analogy between telling a lie and spreading fake news does not demonstrate that fake news must lack truth. On the contrary, we seem to have good reason to believe that fake news can be true and therefore to be open to the possibility that the case of Clinton’s tax fraud amounts to a (genuine) fake news.

As an initial symptom of the fact that this is the correct answer, notice that although the American voters would correctly believe that Hillary Clinton committed tax fraud, intuitively they would not *know* that she did. Their predicament would resemble in important ways the predicament of the protagonists of standard Gettier cases: epistemic agents who have a justified true belief in a proposition, yet fail to have knowledge on account of the accidental way in which they have acquired their belief (Gettier 1963). American voters would acquire a true belief by consuming a story that has been crafted to deceive them. So, even if we concede that they would

⁴ In the opposite direction, Mukerji suggests that we might feel the temptation *not* to look at an accidentally true news as fake news, and endeavours to explain why we should resist *this* temptation. One interesting suggestion, in particular, is that we might feel inclined not to categorize stories like *Russia Report* as fake news because we are under the influence of the *hindsight bias*. This bias inclines one to perceive a given event as more predictable after it has occurred than before. In the case under discussion, this bias would inadvertently incline us to think, contrary to the stipulation, that the authors of the tax fraud fake did know the story was true. So, we would have the intuition that the tax fraud is not fake news because we could not help thinking that the authors of the story as in possession of evidence that the story is true, and so that they published a genuine news after all.

⁵ Jaster & Lanius’s parallel also considers the case of murdering, which we omit for the sake of brevity.

be reasonable in accepting the story, they would thereby come to believe the truth too accidentally to be creditable with knowledge. And this, one might insist, is exactly what we should expect from fake news: that by consuming it one should not be able to acquire knowledge.

If we take the latter suggestion seriously, we are faced with the question of what epistemic property turns the publication of the story about Hillary Clinton's tax fraud into the assertion of fake news, if lack of truth in Jaster and Lanius' sense is not a plausible candidate. A more promising candidate is the fact that the story *is not based on evidence*, in that it has been invented by its authors out of thin air, and not in response to any reasons of which they were aware. Notice, moreover, that the format properties of the publication—its mimicking the publication of reliable news—ensure that it conveys the *false* implicature of being based on evidence. The latter property of an assertion, that of falsely implying that it is based on evidence, is exemplified by all cases of fake news. So, it is natural to contend that a necessary condition for an assertion to be the assertion of fake news is not that it lacks truth in Jasper and Lanius's sense, but that it merely conveys the false implicature of being based on evidence. This leads us to the following possible definition of fake news:

The assertion of P is the assertion of fake news if and only if (i) it potentially addresses a large enough audience, (ii) with the false implicature of being based on appropriate evidence, and (iii) in an untruthful manner, either because the story is published with the intention to deceive or with no concern with P's truth.

Condition (i) recaps our discussion of sociological properties, and requires that the assertion of fake news, if it does not attain the goal of wide circulation, is at least made in a way that does not undercut the chances of achieving this goal. Condition (iii) is Jaster and Lanius's condition that the assertion of fake news lacks truthfulness in their sense. In comparison with other extant proposals, the novelty is condition (ii), which requires that the assertion of a piece of fake news be not made on the basis of supporting evidence, and in such a way as to generate the implicature of being based on such evidence. In relation to this condition, It is worth emphasizing that Jaster & Lanius (*forthcoming*: 12) briefly consider a similar alternative to the definition they propose: they defend their choice of disregarding it on the grounds that, although it would redefine the concept of fake news entirely, it would be extensionally equivalent to the definition it replaces. This, we submit, is clearly untrue. The two definitions would differ precisely on the ground that they would offer an opposite analysis of untruthful reports made in the form of a publication which, like tax fraud, turn out to be true by accident. While Jaster & Lanius commit themselves to denying that they are fake news, our analysis explains why it is reasonable to consider cases of this sort as genuine fake news.

§2 – *How fake news typically propagates: A descriptive analysis*

Before entering the debate about the normative aspects of the spread of fake news through social media, it is helpful to get clear on the *descriptive* features of the phenomena under consideration.

In particular, we shall shed light on two issues relevant for understanding the proliferation of misinformation, namely the roles any user can play regarding the distribution process of a piece of (fake) news on social media, and the causal factors that explain the behaviour of a user depending of the role they are playing. Let's consider these issues in turn.

At first glance, it might seem natural to endorse a bipartite model that distinguishes between *producers* of fake news—that is, someone whose statements respect the aforementioned requirements of an assertion of fake news—and *consumers* of fake news—that is, someone who belongs to the audience reached by the producer's assertion. The consumption of fake news need not result in a belief in its content. It may be consumed for the sole purpose of being entertained, or for the purpose of the fake news being debunked. Although all these kinds of fake news consumers pose interesting questions of their own, we shall concentrate on the immediately most troubling category of fake news consumers that end up endorsing their content. We call them *recipients*.

This bipartite analysis is not surprising, in that it merely applies the traditional distinction between a speaker and a hearer in testimonial exchanges to the domain of fake news sharing. On closer inspection, however, this model is unsatisfying. The fundamental reason is that recipients of fake news are normally not handed their beliefs over directly from their producers. Rather they endorse contents that appear in their, or their friends', news feed. This brings in a third key role, absent in the model supplied by standard testimonial exchanges: the role of fake news *propagators*. A propagator is an individual who shares fake news with their social network friends by re-posting or re-tweeting it. Users in this category distinguish from the other roles in two peculiar ways. Unlike *recipients*, propagators might not believe in the content of the news they share; more importantly, unlike *producers*, they might not share it with the aim of asserting its content. When someone re-posts content, it might be because they take it to be true, but also because they take it to be curious, thought-provoking, outrageous, or funny. As Regina Rini rightly notices, people are generally happy to claim endorsement of a shared news if it turns out to be true, but they are not if doubts about its accuracy arise (E-48)⁶.

From the perspective of fake news producers, the presence of propagators is fundamental for two reasons: first, because propagators contribute to spreading misinformation throughout the networks; second, because it is likely that by re-posting or re-tweeting some fake news about politics or other normatively-laden topics one would not expect them to endorse, they reach out to some users who do not share their values, thereby providing them with a (pro tanto) reason for believing in the fake news. As it will become clear in the remainder of this section, many causal explanations for the spread of fake news trigger users' values and biases.

Before moving to the second main issue of this section, let us point out that the three roles of social media users in fake news epistemology are rather dynamic. It might be the case that a user plays all three roles at different times and, if it seems quite rare that a random individual behaves as a producer of fake news, it is not implausible to suppose that, despite one's

⁶ According to Lazer et al. (2018: 1095), the spread of fake news on social networks relies on a form of *implicit endorsement* that comes with re-posting a content.

genuine intentions, many active users have at least once played the role of a propagator⁷. No matter how likely it is that a single user performs all three roles on their social media activity, a plausible account of these roles should grant the possibility that this happens.

The important question we're going to address now is *how* this happens, that is, which psychological and systemic mechanisms ensure proliferation of detrimental pieces of misinformation. We shall distinguish between *individual* factors, which trigger a single user's proneness to interact with fake news by sharing it and/or getting to believe it, and *social* factors, which operate at a collective level by increasing the chances that a group of people reacts in a particular way to fake news.

At the level of the single user, a first element that favours propagation of fake news is the low amount of attention users devote to social media contents, in partial response to the overwhelming amount of information that is ready to use online or directly flows in each one's news feed. As a recent research report by Microsoft Canada has shown, digital lifestyles decrease the level of sustained attention we deploy in our activities, thereby reducing our ability to maintain prolonged focus during repetitive activity.⁸ Considerations about the level of attention we devote to social media content help us explain the spread of fake news at a more general level. In their recent study (2017), Qiu and colleagues demonstrated that "both information load and limited attention lead to low discriminative power" (5), that is, to our inability to assess whether the new contents we see appearing in our news feed are reliable. However, it is not clear as of today whether low discriminative power straightforwardly leads to virality of misinformation: in particular, Qiu and colleagues have recently retracted the original result of their study by admitting that they have not found evidence that "low-quality information is just as likely to go viral" than high-quality information is (2019).

Regardless of whether low levels of sustained attention contribute to the spread of fake news, it is a fact that falsehoods—at least on Twitter—diffuse "significantly farther, faster, deeper and more broadly than the truth in all categories of information" (Vosoughy et al. 2018: 1147): in a word, the study of Vosoughy and colleagues shows that falsehoods are "70% more likely to be retweeted than the truth" (1149). Furthermore, those users who are more likely to get trapped in cascades of fake news have, in general, less followers than those who spread the truth; follow fewer people; are less active on the social network; are verified less often; and have been on social media for less time. To our best knowledge, there is no such study that provides a comprehensive explanation of why those users are more prone to share fake news than others. We do have evidence, though, that further factors contribute to accounting for our general propensity to believe in—and share—misinformation.

⁷ As *Statista.com* reports, 23% of participants to a 2016 U.S. survey admitted having either shared a political news story that it turned out to be made up or shared a political news story that they knew at the time it was made up (see <https://www.statista.com/statistics/657111/fake-news-sharing-online/>). It is thus obvious to expect a higher percentage of individuals who either admitted having shared a political fake news or did not realize (or admit) they have shared a political fake news.

⁸ Microsoft Attention Span Research Report, Microsoft Canada (accessed via Scribd at <https://www.scribd.com/document/265348695/Microsoft-Attention-Spans-Research-Report>).

One of them is cognitive biases, in particular confirmation bias and desirability bias. The former leads a user to seek out information that confirms what they already believe and disregards what clashes with one's opinion. As Tom Nichols has interestingly pointed out, confirmation bias is somewhat a survival tool (2017: 64), which activates to prevent us from getting crushed by the flow of information on social media. The problem, though, is that this bias makes us insensitive to evidence against our beliefs, preconceptions, and normative views, thereby compromising our rational capacity to evaluate new information. For example, if a climate-change denialist is presented with evidence that the Earth is warming, they will be less attentive to, or more prone to disregard it.

The latter bias, i.e. desirability bias, leads a user to assign more credibility to desirable pieces of information than to undesirable ones, thereby displaying a tendency to "update their prior beliefs to incorporate new and confirming information more than new but disconfirming information" (Tappin et al. 2017: 1143). In the case of fake news, desirability bias might operate as follows: for example, if one is inclined to believe that an increase in a city's crime rate is due to the presence of illegal immigrants, then one will be more prone to believe in fake news showing and condemning the arrival of hundreds of refugees.⁹ Similarly, if one is already convinced that Hillary Clinton is an immoral and perverted individual, then one will be prone to believe in a fake story about her involvement in child sex trafficking.

As we have tried to point out, cognitive biases constitute a relevant factor to account for the spread of fake news on social media, namely one that relies on shortcomings of human rationality before a piece of information. Other—possibly more surprising—factors include what psychologists call *memory-based mechanisms* and *fluency effects*. Both activate in cases in which one acknowledges that some social media content is fake news and updates their beliefs accordingly. As a result of the former, several bad things may happen: first, one may fail to recall whether some content is part of the fake news or of a reliable piece of information; second, one may mistakenly take a claim to be issued by a reliable source when it is in fact part of the fake news; even worse, one may mistakenly attribute such claim to themselves, thereby self-ascribing flawed evidence and making their belief resistant to disconfirmation (Levy 2017: 29).

The latter, instead, explain why repeated exposure to fake news story which a user knows to be fake increases the chances that they forget its source, thereby leading them to misattribute the originator to a reliable source and hence to have reasons for believing in the fake news story (30). Fluency effects are particularly relevant as a causal factor for explaining proliferation of fake news in social media: as it will become clear in the rest of this section, several structural features of social networks facilitate repeated appearance of contents in one's news feed and hence enhance the likelihood that a user's fluency of processing is affected over time.

Let us then consider three relevant social explanations for the spread of fake news on social media, namely epistemic bubbles, echo chambers, and group polarization. A recent paper by Thi Nguyen (2018) points out that both epistemic bubbles and echo chambers amount to

⁹ Detailed analysis of fake news targeting refugees is available at <https://teyit.org/en/how-is-false-information-used-worldwide-to-target-refugees/>.

structures of exclusion, in that they prevent large groups of social media users from getting aware of—or from taking in due account—some kinds of information. Thus, both contribute to generating and sustaining an ideological detachment, insofar as one ends up lacking due contact with—or being prevented from appreciating the relevance of—information and opinions that contrast with the view of the groups one belongs to.

However, epistemic bubbles generate and operate in a different way than echo chambers. Epistemic bubbles amount to social structures that impede distribution of a complete range of information by omitting relevant testimony from sources endorsing a rival perspective. This process of exclusion is supported by tendencies at an individual level such as confirmation bias or selective exposure, and may generate what Nguyen calls *bootstrapped corroboration*, namely the tendency to overestimate one's self-confidence in epistemic judgments based on wide agreement one is likely to find among fellows within an epistemic bubble (4). However, it's in the nature of bubbles that they can be easily popped-up: specifically, epistemic bubbles disappear insofar as some members are exposed to excluded information.

Echo chambers obstruct consumption of information more actively: in particular, membership within an echo chamber requires that its members preliminary accept a set of meta-beliefs that distribute epistemic credit asymmetrically between insiders and outsiders, by overinflating the epistemic trustworthiness of the former and by overdeflating that of the latter irrespective of their actual epistemic worth. Epistemic discredit can be cashed out in terms of epistemic unreliability, but also in terms of intellectual vices such as epistemic maliciousness, close-mindedness, or dishonesty.

The expected counterpart of bootstrapped corroboration for epistemic bubbles is what Nguyen calls a *disagreement-reinforcing mechanism*, which pre-emptively cancels out the effect of counter-evidence and conflicting information by alerting members of an echo chamber that this is what they should expect from epistemically corrupt outsiders. By contrast, the mere fact that insiders share similar beliefs and assess information in the same way provides each of them with further evidence that their fellows are trustworthy, thereby generating boosts of the inner levels of credence within the echo chamber. This mechanism highlights the difference between epistemic bubbles and echo chambers: paradoxically enough, in the latter, unlike the former, wider exposure is likely to backfire, that is, to contribute to discrediting contrary opinions and increasing inner trust among the members of an echo chamber.

Despite their differences, it should be evident that both epistemic bubbles and echo chambers foster the spread of fake news on social media. For example, some false posts containing claims about the arrival of new refugees on the Italian coasts, associated with unrelated pictures of ships full of foreign migrants, are likely to spread within nationalistic networks that condemn policies of asylum to refugees. If members of such groups tend to disregard fact checking sources about the truth of the posts in question yet are ready to stop believing in their contents as soon as counter-evidence sneaks in the community, the filter bubble can easily pop-up. The remaining worry at that point would just be that reiteration of such posts may generate the aforementioned fluency effects. In contrast, if this story involves a group in

which fact checking sources are dismissed as corrupt agencies controlled by leftist parties, then not only people will be resistant to the evidence that the posts were fake, but they will consider such information as the obvious reaction of leftist conspirators, thereby strengthening their belief that multitudes of illegal immigrants are docking at Italian ports, as the analysis of echo chambers predicts.

The last factor we shall consider in this section is group polarization, namely a group's tendency to display more extreme beliefs and attitudes than the ones its members possess when taken as single individuals. It makes sense to end our list of factors explaining the spread of fake news over social networks with group polarization because such tendency gathers many—if not all the—*aforementioned* elements. In particular, low levels of attention devoted to social media contents explain why people are reluctant to change opinion (Riva 2017; Sunstein 2017), while confirmation bias and echo chambers explain why collective dynamics lead members of a group to polarize toward extreme views by increasing self-confidence in their opinions and taking wide (dis)agreement with insiders (outsiders) as confirming evidence that they are right (Nguyen 2018; Sunstein 2017)¹⁰. In other words, group polarization intensifies the effects of the other factors, thereby increasing the chances that users accept and share fake news whose content supports their view, and dismiss as conspiracy theory any information that contradicts their opinions.

Now, it should be stressed that the proposed considerations merely provide a concise analysis of the most recent works in a field that is growing fast. On the one hand, this justifies why we should refrain from drawing strong conclusions about which factors best account for the descriptive causes of the spread of misinformation on social networks. On the other, it should also suggest that we will hopefully be in a position to have more conclusive answers to such a question in the upcoming years.

§3 – *Normative questions*

In this final section we turn to the normative dimension of fake news. In the last section we have proposed a tripartite model, which distinguishes between three roles: the *producer* of fake news, its *propagator*, and its final *recipient*. In accordance with this model, it is natural to address the normative dimension of fake news by addressing the normative dimension of each of the *aforementioned* roles. We shall do so with two important limitations. First, we shall disregard the *moral* dimension of fake news, such as the moral shortcomings of the propagation of misleading contents. Instead, we shall merely focus on the *epistemic* dimension of fake news. Secondly, although we will say something about the epistemology of fake news *propagation*, we shall primarily focus on the epistemic dimension of fake news *reception* and entirely disregard the epistemic aspects connected with fake news *production*. This second limitation is due to the fact that fake news producers are primarily pursuing a political or financial agenda. Fake news propagators and recipients, on the other hand, can be regarded as primarily engaging in an *epistemic* enterprise, that of gathering and transmitting information about the world. So, we regard the epistemic assessment of this practice as entirely appropriate and as potentially performing a

¹⁰ For a comprehensive analysis of exposure to ideologically diverse news on social media, see Bakshy et al. (2015).

corrective role. If the way in which fake news recipients (and propagators) engage in their practice turns out to be epistemically defective, they should be required to modify their behaviour accordingly.

In response to the question about the epistemic merits of fake news reception, three possible answers stand out in the literature. According to one first answer, defended by Rini (2017), fake news recipients are ordinarily acting in an epistemically *virtuous manner*. Rini does not think the widespread diffusion of fake news is for this reason a good thing. However, she believes that the problem is not located at the individual but at a systemic level. A second answer shares the idea that the problem of fake news is primarily located at the systemic level, but it individuates the reason in the fact that fake news recipients are ordinarily epistemically *blameless* (Nguyen 2018). Epistemic blamelessness is normally taken to be necessary for epistemic virtue or for the possession of some other epistemically valuable standing such as justification. So, acceptance of the first answer commits to accepting the second answer too (in fact Rini, and perhaps Nguyen, accepts both). According to a third answer, fake news recipients are epistemically defective and manage their beliefs out of epistemic vice (Cassam 2016). Proponents of this answer believe that the problem of fake news can be primarily resolved by promoting a reform of individual epistemic behaviours.

Both Rini's and Nguyen's answers presuppose that fake news consumption happens in specific environments—though for different reasons. According to Rini the reception of fake news can be described in the relevant respects as resulting from a specific kind of *testimonial exchange*, where a propagator typically shares fake news over a social medium and a recipient comes to believe its content because they read it in their or their friends' news feed. According to Rini, a fundamental fact about news consumption on social media is that the news propagators that most likely appear in one's news feed usually belong to one's network of friends and contacts. These friends and contacts are normally selected because they share basic commitments to fundamental moral and political values. This, on Rini's view, justifies the presumption that the contents they share—at least when they pertain to normative domains—are worth being accepted. For given the way they have been selected, it is rational to presume that, at least by the lights of one's own standards, they tend “to get normative questions right” (2017: E-51). This, for Rini, vindicates as epistemically virtuous the practice of uncritically accepting normative news shared by one's contacts on social media (Rini speaks of the virtue of *epistemic partisanship*). This is true also when the news is fake, provided that it has the right subject matter. Here we consider two possible cases. Although we believe Rini may be right about the first, we contend that she is wrong about the second.

The first kind of news that we may reasonably accept in a partisan way from our peers is straightforward normative claims, namely claims that openly reflect one's allegiance to some set of basic values—for instance, claims to the effect that a given policy would be unjust, morally objectionable, etc. Rini's suggestion is probably right in this case: insofar as I have reasons to believe that my contacts share my value commitments, *other things being equal* I am justified to trust their normative claims (or other parties' normative claims, if my contacts agree with them) more

than I am justified in trusting the normative claims made by someone who is not in my partisan network. This argument does not exclude the possibility that one be fooled into believing fake news if a co-partisan fellow shares one: rather, it explains why there might be nothing wrong in the recipient when they trust a partisan fellow who propagates fake news about normative claims.

However, fake news rarely pertains to normative claims in this sense. It more likely pertains to what Rini calls *normatively relevant claims*, namely descriptive claims that have a clear bearing on some normative question. Claims of this sort include, for instance, stories about the actions of a politician, the likely consequences of a given policy, etc. Rini contends that it is epistemically virtuous to trust in a partisan way one's contacts also when they share news attesting to some normatively relevant claim. Here Rini's reasons are less than conclusive.

Suppose that one of our contacts shares a news report, according to which P is true. When P is normatively relevant, according to Rini, the fact that our contact decided to share the report provides us with an epistemic reason for accepting P. This is so, according to her, because the decision to share the news report was a *normative* decision, reflecting our contact's assessment of the *importance* of the question about whether P is true (2017: E-52). We agree that the fact that the news report has been shared by one of our contacts makes it rational to expect that the news report will be *important* to us as well. However, the question is not about whether we're rational in deciding to spend some time in consuming the relevant news report. The question is about whether we are rational in accepting what the news report says, namely P. And while our shared commitments with our contact sustain an affirmative answer to the first question, there's no reason to suppose that they sustain an affirmative answer to the second as well. The accuracy of the news report, differently than the importance that we attach to it, is not a matter of agreement between our values and the values of our contact; it is a matter of the report's agreement with the facts.

As said, also Nguyen (2018) addresses the question about the epistemic worthiness of fake news consumption and reception against the background of his analysis of the environments over which fake news tend to proliferate. As we have seen in the last section, for Nguyen these environments typically exhibit the features of an echo chamber, which obstructs one's members' free access to information not necessarily by omitting or excluding certain specific sources, but by actively discrediting them. This analysis has important consequences with respect to the question of whether members of an echo chamber can be criticized, from an epistemic point of view, for the way they manage their beliefs. Nguyen explicitly suggests that members of an echo chamber, despite complying with the unrealistic distribution of epistemic credit required by the relevant meta-beliefs, can act "much like a reasonable epistemic agent" (15). Nguyen goes so far as to suggest that being "trapped" within an echo chamber does not necessarily prevent an epistemic agent from being epistemically virtuous (15).

In order to investigate Nguyen's suggestion in more detail, and to assess its implications for the normative status of fake news reception, it is useful to consider the following example¹¹. Suppose that Oliver has been raised within an echo chamber according to which the collapse of

¹¹ The example is Nguyen's revised version of a case originally proposed in Cassam (2014).

the twin towers was an inside job. His parents and all sorts of epistemic authorities within his community have continuously reinforced the beliefs of the echo chamber. These beliefs, among other things, require distrusting the official story about 9/11, as a conspiracy theory that has been created and widely promoted in the attempt to hide the real explanation for the towers' collapse. One day a website X, which Oliver takes to be a reliable source¹², publishes an (otherwise incredible) story—we might suppose it meets all the aforementioned conditions for being fake news—which, if true, would confirm the echo chamber's core beliefs. Oliver deliberates about what he should do with the story, and after careful reflection he resolves to believe it. How should we assess Oliver's epistemic conduct?

To begin with, notice that Oliver seems to have acquired the echo chamber beliefs in a reasonable manner, namely by trusting the epistemic authorities of his epistemic community. So, he is not to be blamed for being a member of the echo chamber. We can also imagine that Oliver has exercised many epistemic virtues in arriving at his belief. He has been active in seeking new sources of information, before he stumbled in the story published by X; he has thoroughly investigated his sources and has been scrupulous in assessing their credibility in light of his background meta-beliefs, etc. Finally, by accepting X's story he has resolved to believe in a way that accords with his reasoned assessment of the evidence at his disposal. According to Nguyen's analysis, all this shows that Oliver is epistemically blameless for believing in the fake news. His epistemic agency is epistemically virtuous overall; it's the community's epistemic structure that is epistemically vicious, and that explains why, despite his reasonable behaviour, Oliver fails to achieve an epistemic good.

At this point, one might suspect Nguyen's analysis perfectly aligns with Rini's diagnosis of fake news recipients. On closer inspection, two important details show that Nguyen endorses a more reasonable thesis than Rini does. First, Nguyen's claim is *weaker*. According to him, one can be epistemically virtuous *despite* being the recipient of fake news, insofar as they undertake evidence-gathering processes that display all sorts of traditional epistemic virtues. In contrast, Rini contends that one can be epistemically virtuous *qua* recipient of fake news by displaying the virtue of epistemic partisanship, that is, by assigning greater credibility to a testifier—e.g. a fake news propagator—because they know they share a normative affiliation with them. Second, Nguyen's thesis is *narrower*, in that he limits the possibility of attributing epistemic virtue to fake news recipients who are blameless for having been raised in echo chambers. In contrast, Rini is willing to grant epistemic partisanship to any fake news recipient, irrespective of the social and epistemic features of their community, to the extent that they trust a propagator based on the fact that they have analogous or similar normative values.

So far, we have provided reasons for rejecting the strong thesis that fake news recipients undertaking epistemically partisan assessments may be virtuous. Yet we have argued in favour of a weaker claim, namely that fake news recipients raised in echo chambers may be blameless for believing in fake news and epistemically virtuous in their methods of gathering new information.

¹² Let us stipulate that Oliver's belief about the reliability of website X amounts to one of the meta-beliefs he inherits from the echo chamber he is part of.

Were it to be the case that all recipients of fake news are members of echo chambers by default, the latter thesis would settle the score. However, it seems plausible to concede that fake news may propagate in epistemic bubbles and less vicious communities too. In this scenario, Cassam’s considerations about the individual responsibility of fake news recipients play a fundamental role, or so we shall argue.

Consider a revised version of Oliver’s case, in which the subject is not a member of an echo chamber: rather, he simply read a piece of fake news according to which the towers’ collapse was due to explosives planted in the buildings by government agents; he got interested in the conspiracy theory; and he found further evidence supporting the thesis that aircraft impacts could not have caused the towers to collapse, and other stories of this sort. As a result, Oliver convinces himself that the official explanation of 9/11 has been created to cover up the real causes of the collapse.

According to Cassam (2016), Oliver is blameworthy, as he fails to comply with the norms of responsible inquiry, which at least require having a good sense of when one is in danger of being deceived (163). The relevant defects of his epistemic agency can be cashed out in terms of intellectual vices: Oliver displays close-mindedness, as he fails to take into adequate consideration counter-evidence against his theory; “prejudicial dysfunction” (Fricker 2012: 340), as he misplaces trust in other people by granting excessive credibility to wrong sources and denying appropriate credibility to legitimate sources; and lack of thoroughness, as he fails to acknowledge that alleged evidence in support of the conspiracy theory does not stand up to scrutiny.

Two short remarks are in order. First, notice that the relevant set of intellectual vices that play a role in the domain of fake news epistemology is much broader than the items included in Oliver’s case. To make just one example, it seems plausible to contend that an epistemic subject who behaves as Rini describes in the case of normatively relevant claims—i.e. someone who systematically attributes excessive credibility to co-partisan fellows based on shared normative values in situations in which one is supposed to seek *factual* rather than *normative* evidence—displays a peculiar epistemic vice, which we shall call *partisan gullibility*. Furthermore, our analysis has mainly focused on the epistemic attitudes of fake news recipients, yet considerations of a similar sort may be—in fact, should be—offered about fake news producers and propagators, who will likely reveal to possess other intellectual vices¹³.

Second, it is worth pointing out that being responsible inquirers, thus failing to display and cultivate intellectual vices, is not only key to our intellectual flourishing qua single social media users, but also to the epistemic aspirations of our epistemic communities. As the literature about fake news proliferation shows, the beliefs we form through social networks and our actions qua users are likely to affect not only our own epistemic assessments, but also those of other community members.

One might suspect that the proposed analysis of the normative aspects of fake news propagation leaves us with two different options, namely one relieving single recipients of

¹³ Appeal to a virtue epistemology as a relevant strategy to address the spread of fake news has been recently proposed by Heersmink (2018) and Smart (2018).

responsibility for believing in fake news and another blaming them for their irresponsible and vicious attitudes. In fact, the two options are compatible with each other, as they address different scenarios: namely, one in which the spread of misinformation throughout a network is due to peculiar features of its social and epistemic structure (such as echo chambers); another in which single users' epistemic defects contribute a great deal to the pervasiveness of fake news.

As things stand, it seems reasonable to look with suspicion at any attempt to provide correctives for the spread of fake news that fail to take into consideration both systemic and individual responsibility¹⁴. Conversely, we hope that our analysis has contributed to motivating the need for comprehensive and nuanced solutions to a problem that threaten to hinder the epistemic progress of the entire global village.

Conclusion

This paper surveyed the growing literature on fake news in the attempt to address three fundamental issues for fake news epistemology. First, we analysed the debate about what fake news is and provided a novel account that purports to avoid some problems affecting other extant views on the market. Then, we inquired into the descriptive causes of fake news proliferation to shed light on psychological relevant features of social media users as well as on social, i.e. structural, dynamics that social media sharing gives rise to and sustains. Finally, we discussed some normative issues related to fake news reception: in particular, we argued against 'all-or-nothing' views that either blame or discharge individual users from any responsibility for their attitudes as consumers of news on social networks. In contrast, we defended the thesis that individual users may be blameworthy for believing in fake news unless they have been raised in echo chambers and never had an opportunity to acknowledge the epistemic partiality of their epistemic environment.

Bibliography

- Allcott, H. and Gentzkow, M. (2017). Social Media and Fake News in the 2016 Election, *National Bureau of Economic Research*, Working Paper 23089.
- Bakshy, E., Messing, S., and Adamic, L. A. (2015). Exposure to Ideologically Diverse News and Opinion on Facebook. *Science*, 348 (6239), pp. 1130-1132.
- Cassam, Q. (2016). Vice Epistemology. *The Monist*, 99 (2), pp. 159–180.
- Chisholm, R. M., and Feehan, T. D. (1977). The Intent to Deceive. *Journal of Philosophy*, 74, pp. 143–159.
- Dentith, M. R. X. (2017). The Problem of Fake News. *Public Reason* 8 (1-2), pp. 65–79.
- Frankfurt, H. G. (2005). *On Bullshit*. Princeton: Princeton University Press.
- Gelfert, A. (2018). Fake News: A Definition. *Informal Logic*, 38 (1), pp. 84–117.
- Gettier, E. (1963). Is Justified True Belief Knowledge?. *Analysis*, 23 (6), pp. 121–123.

¹⁴ Notice that the solutions proposed by Rini (2017) and Nguyen (2018) seem to be affected by this limitation. In this sense, our diagnosis falls more in line with Lazer et al. (2018).

- Fricker, M. (2012). Epistemic Injustice and Role for Virtue in the Politics of Knowing.” In J. Greco and J. Turri (eds.), *Virtue Epistemology: Contemporary Readings*, Cambridge (MA): MIT Press, pp. 329-350.
- Habgood-Coote, J. (forthcoming). Stop Talking About Fake News. *Inquiry*.
- Heersmink, R. (2018). A Virtue Epistemology of the Internet: Search Engines, Intellectual Virtues, and Education. *Social Epistemology*, 32 (1), pp. 1-12.
- Isenberg, A. (1973). Deontology and the Ethics of Lying. *Aesthetics and Theory of Criticism: Selected Essays of Arnold Isenberg*, Chicago: University of Chicago Press, pp. 245–264.
- Jaster, R., and Lanus, D., (2018). What is Fake News?. *Versus* 2 (127), pp. 207–227.
- Klein, D., and Wueller, J. (2017). Fake News: A Legal Perspective. *Journal of Internet Law*, 20 (10).
- Lackey, J. (2013). Lies and Deception: An Unhappy Divorce. *Analysis*, 73, pp. 236–248.
- Lazer, D. M., Baum, M. A., Benkler, Y., Berinski, A. J., Greenhill, K. M., Menczer, F., ... and Schudson, M. (2018). The Science of Fake News. *Science*, 359 (6380), pp. 1094-1096.
- Levy, N. (2017). The Bad News about Fake News. *Social Epistemology Review and Reply Collective*, 6 (8), pp. 20-36.
- Mahon, J. E. (2008). Two Definitions of Lying. *International Journal of Applied Philosophy*, 22, pp. 211–230.
- Mukerji, N. S. (forthcoming). What is Fake News?. *Ergo: An Open Access Journal of Philosophy*, 18.
- Nichols, T. (2018). *La conoscenza e i suoi nemici*. Rome: Luiss University Press.
- Nguyen C. T. (2018). Echo Chambers and Epistemic Bubbles. *Episteme*.
- Primoratz, I. (1984). Lying and the “Methods of Ethics”. *International Studies in Philosophy*, 16, pp. 35–57.
- Qiu, X., Oliveira, D. F., Shirazi, A. S., Flammini, A., and Menczer, F. (2017). Limited Individual Attention and Online Virality of Low-Quality Information. *Nature Human Behaviour*, 1 (7), 0132.
- Qiu, X., Oliveira, D. F., Shirazi, A. S., Flammini, A., and Menczer, F. (2019). Retraction Note: Limited Individual Attention and Online Virality of Low-Quality Information. *Nature Human Behaviour*, doi: 10.1038/s41562-018-0507-0.
- Rini, R. (2017). Fake News and Partisan Epistemology. *Kennedy Institute of Ethics Journal* 27 (S2), pp. E43–E64.
- Riva, G. (2018). Fake News. Bologna: Il Mulino.
- Smart, P. (2018). (Fake?) News Alert: Intellectual Virtues Required for Online Knowledge! *Social Epistemology Review and Reply Collective*, 7 (2), pp. 45-55.
- Sunstein, C. R. (2017). #republic. *La democrazia nell'era dei social media*. Bologna: Il Mulino.
- Tappin, B. M., van der Leer, L., and McKay, R. T. (2017). The Heart Trumps The Head: Desirability Bias in Political Belief Revision. *Journal of Experimental Psychology: General*, 146 (8), 1143.
- Vosoughi, S., Roy, D., and Aral, S. (2018). The Spread of True and False News Online. *Science*, 359 (6380), pp. 1146-1151.

Williams, B. (2002). *Truth and Truthfulness: An Essay in Genealogy*, Princeton: Princeton University Press.