MISINFORMATION AND INTENTIONAL DECEPTION:

A NOVEL ACCOUNT OF FAKE NEWS

Forthcoming in N. Snow & M. S. Vaccarezza (eds.), *Virtues, Democracy, and Online Media: Ethical and Epistemic Issues*, Routledge.

MICHEL CROCE

University of Edinburgh & University of Lisbon

TOMMASO PIAZZA

University of Pavia

ABSTRACT. This chapter introduces a novel account of fake news and explains how it differs from other definitions on the market. The account locates the fakeness of an alleged news report in two main aspects related to its production, namely that its creators do not think to have sufficient evidence in favor of what they divulge and they fail to display the appropriate attitude towards the truth of the information they share. A key feature of our analysis is that it does not require that fake news must be circulated with the intention to deceive one's audience. In this way, our account overcomes a potential limitation of the current philosophical discussion about fake news, which appears to individuate the main problem with this phenomenon in the fact that fake news consumers are misled and misinformed. In contrast, the proposed analysis shows that an additional (and perhaps equally fundamental) problem uncovered by the spread of fake news is a widespread pathological relationship with information, one on which we consume information not to satisfy our interest in the truth but to strengthen our social identities and quench our hunger for social recognition.

KEYWORDS: fake news; intention to deceive; clickbait; misinformation; applied epistemology

[penultimate version]

1. Introduction

In the last few years, philosophers have started to inquire into the phenomenon of *fake news*: what it is, why social media prove particularly hospitable to its proliferation, who should be held responsible for its detrimental effects, and what could be done to counteract such effects. Among the reasons why philosophy can help us get clear on these issues, a legitimate role is played by concerns with our epistemic welfare and the overall wellbeing of our communities.[1] It should come as no surprise that the current spread of fake news ruins the quality of our epistemic environment.

---

[1] Habgood-Coote (2019) disputes the thesis that there is something philosophically relevant to discuss about the term *fake news*. For a reply to his view, see Pepp, Michaelson, Sterken (2019a).

But considering the spread of fake news merely as an epistemic problem is overly simplistic. If the quality of the goods we exchange in the epistemic market is poor, our decision-making processes rely on an impoverished factual basis and, in turn, this threatens the wellbeing of our communities on many levels: notably, it worsens the welfare of our democratic institutions and the citizens' health, and it increases the negative effects of stereotypes and cognitive biases.

These considerations rightly elicit a concern with the responsibility for producing and sharing fake news as well as with the measures that can be adopted to counteract its spread. We shall say something about the former in the next section, but we tackle both extensively elsewhere (Croce and Piazza 2020). This chapter has a more fundamental aim, namely to offer and motivate a novel account of fake news (section 2) that takes inspiration from extant views on the market but purports to avoid two problems that affect the current philosophical discussion (section 3).[2]

The first problem pertains to the intentions of fake news producers. Duncan Pritchard (*forthcoming*) understands the propagation of fake news as a form of a testimonial exchange in an unfriendly context replete with bad epistemic actors who intend to deceive their audience. On our view, a deceitful intent is not a necessary feature of fake news: a more plausible story of how fake news is fabricated and propagates has it that some producers—in particular, clickbait factories—craft minimally plausible stories to be clicked on and present them as news merely with the interest to earn money out of the clicks that their contents generate. We shall argue that, in this scenario, fake news producers need not make a big effort in creating stories that will be widely believed by their consumers, so long as these stories satisfy non-epistemic desiderata such as being engaging and pleasing to their intended audience.

The second issue we aim to address concerns the reasons why social media users consume fake news. We argue that the current philosophical discussion unduly restricts the focus of this phenomenon to the epistemic factors that lead people to consume fake news, namely a misplaced hunger for reliable information. It is surely the case that many people consume and share fake news because they believe in it. But a more appropriate account has it that fake news consumption is (at least partly) fostered by a lack of concern for the truth—and other epistemically relevant factors—of the shared contents. More precisely, we attempt to show that fake news can be consumed as a way to satisfy a range of non-epistemic needs having to do with one's socio-political convictions, emotional bonding as a member of a community, and need to be entertained—in a word: a hunger for social recognition. We will conclude that this pathological relation with information is a second major problem connected with the circulation of fake news, surprisingly neglected in the current literature, which our account helps focussing upon.

---

[2] See Piazza and Croce (2019) for an initial sketch of this account.

2. *Fake News and Its Properties*

A promising strategy to arrive at a definition of fake news is to individuate and discuss the properties that a piece of information has to display to count as an instance of fake news.[3] Two preliminary remarks are in order.

The first is methodological in spirit. Getting clear on how to understand the concept of fake news is particularly important because the common usage of this term is far from stable. (Think, for example, of how politicians weaponize the term 'fake news' as a way to label and discredit political opponents or attack the media.)[4] While we aim at offering a definition of fake news that aligns with its non-weaponized usage in everyday contexts, our account purports to clarify the meaning of this notion in such a way that it becomes easier to identify and counteract the detrimental effects of the spread of misinformation. Thus, on our approach, conceptual analysis is at the service of an ameliorative project.

The second remark has to do with the genre fake news should be taken to belong as a species. Most commentators agree that fake news is propositional content which can be conveyed in a variety of ways—e.g., by articles, pictures, videos, and the like (Gelfert 2018: 103-104; Jaster and Lanius 2018: 208; Mukerji 2018: 930; Rini 2017: E-44). For the sake of simplicity, in this chapter we will assume that fake news is the propositional content asserted—namely stated or implied—in any of the ways just mentioned.

Let us clarify this a bit further. That fake news often is the literal content of an assertion should come as no surprise: for example, almost all versions of the posts or fake news reports concerning the famous Pizzagate conspiracy literally stated that Hillary Clinton was involved in a child trafficking ring. However, it seems plausible that, in some cases, fake news hides in the implicated content of an assertion. By reporting the number of burglaries that occurred in town after the arrival of the refugees, a newspaper would not literally say, but could certainly imply, that the refugees are responsible for the burglaries.[5] If this were false or ungrounded, it would be natural to classify the initial report as fake news, although what it states is true.[6]

---

[3] We have offered a more extended analysis of these properties in Piazza and Croce (2019).

[4] In particular, the President of the United States, Donald Trump, got to the point of personalizing the referent of this term by saying to a CNN reporter "You're fake news!" or labeling media sources that shed light on the shortcomings in his administration "The Fake News Media" (https://cpj.org/reports/2020/04/trump-media-attacks-credibility-leaks.php.). Further evidence of how politicians weaponize the term 'fake news' can be found here: https://theconversation.com/the-real-news-on-fake-news-politicians-use-it-to-discredit-media-and-journalists-need-to-fight-back-123907.

[5] A more detailed version of the example is offered in Jaster and Lanius (2018).

[6] Mukerji (2018) 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.

We can now ask what properties turn an assertion in our broad sense into an assertion whose literal or implied content is fake news. We shall focus our attention on four kinds of properties—namely, epistemic, intentional, format, and sociological properties: for the sake of simplicity, we inspect them in reverse order, starting from the properties that require less discussion.

Sociological properties have to do with the upshot of a piece of fake news in terms of circulation. Some theorists (Gelfert 2018; Pepp, Michaelson, and Sterken 2019b) contend that a piece of fake news, to be so, must attain the goal of wide circulation and uptake. Rini (2018) merely requires that fake news be transmitted with *the goal* of wide circulation and re-transmission. On our view, wide circulation is *ceteris paribus* sufficient for a message to convey fake news, but 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. 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.

These considerations have already called attention to format properties, namely to the mode in which some propositional content has to be presented for it to constitute a piece of fake news. There seems to be wide consensus over the idea that fake news, to be so, must be *presented* as news, where this means that it must possess certain format features that make it appear like any other piece of real news (Gelfert 2018: 103; Pritchard *forthcoming*). For the purposes of this chapter, we shall remain pluralists about *where* a piece of fake news has to be displayed—that is, whether fake news requires online publication (Klein and Wueller 2017: 6) and what relationship it has with traditional media[7]—insofar as such content is distributed over a medium that is suitable for wide circulation.

Let us now inquire into the two fundamental, albeit more problematic, properties of fake news. With the term *intentional properties*, we want to target the *intentions* that motivate the act of making a given assertion and the *attitude* with which one makes it. At first glance, it would seem that to assert a piece of fake news one has to have the *intention* of deceiving their audience (Dentith 2017: 66; Gelfert 2018: 108; Pritchard *forthcoming*). After all, genuinely mistaken reports which *inadvertently* offer an incorrect picture of some state of affairs are not fake news precisely because they are meant to inform, rather than to instill false beliefs in, their audience.

Yet, things are more complicated than it initially appears, in that there are also cases of fake news whose producers would appear to lack the intention to deceive their audience. The usual example is that of the Macedonian clickbait factories, multiple websites managed by teenagers (most of which from the small town of Veles) filled with appealing and sensationalist fake news on

---

[7] For further considerations on this point, see, e.g., Gelfert (2018), Mukerji (2018), and Rini (2017).

a daily basis for the sole purpose of generating traffic and of gaining money proportionally from the adverstising. Reportedly, this industry was very profitable.[8] As we shall see more in detail in the next section, it is to a certain extent controversial how this example should be interpreted. For the time being, suffice it to say that clickbaiters—like bullshitters, in general—do not 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 sole reason that it suits their practical ends. At this stage, we can thus register that the attitude of clickbaiters *might not* be one of intending or attempting to deceive their audience, but rather one of indifference toward the truth of the contents they share. As a consequence, a necessary (intention-related) condition for something to be fake news should take a disjunctive form: it must be *either* asserted with the intention to deceive or with a lack of concern for its truth.

Although it points in the right direction, the latter claim must, however, be qualified by making explicit *whose* attitudes matter to the status of fake news. Suppose one interprets the act whereby a propagator re-posts a content on social media as an assertion. If also the attitudes of the propagators mattered, in many cases we would have to deny that what circulates widely is fake news. For one reason why fake news circulates widely is arguably that a portion of those who contribute to its propagation sincerely believe what they say. This problem is easily pre-empted by making explicit that the disjunctive condition above pertains to the assertions, and the intentions underwriting them, of the *initial* creators—and not of the many propagators—of fake news. Thus, throughout the remainder of this paper, we shall take for granted that the assertions which our account refers to are the assertions of their initial producers. Towards the end of this section, we shall address a related complication.

Considerations about truth and falsity of one's assertions bring us to the final property of fake news, which we call *epistemic. Falsity* is a natural culprit when, in this broader sense, it comes to the epistemic properties of the assertions whose content is fake news (Rini 2017: E-45). As we have already seen, however, one may achieve the goal of disinforming one's audience also by making a *misleading* assertion, namely by generating a false implicature through the statement of a true content (Allcott and Gentzkow 2017; Gelfert 2018). With the words of Jaster and Lanius, we could simply put it that fake news is characterized by a *lack of truth*, where this can "come about through the falsity of a news report or … through its misleadingness" (2018: 210).

An intriguing question is whether a lack of truth is a *necessary* condition for an assertion to be the assertion of fake news. Suppose *Russia Today* fabricates the story that Hilary Clinton committed tax fraud (Jaster and Lanius 2018: 218). 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

---

[8] For a detailed description see, e.g., https://www.wired.com/2017/02/veles-macedonia-fake-news/.

considerably damaged by the allegation that she didn't pay her taxes. So, despite 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 and Lanius maintain that we might feel tempted to look at the *Russia Today* case as involving a piece of fake news, but argue that we should not give in to this temptation.[9] Their argument, however, can be easily resisted. According to them, we should distinguish between successfully lying and merely attempting to lie, which is what we do when we inadvertently assert a truth we do not believe in with the intention to deceive. In the same way, we should distinguish between spreading fake news and merely attempting to do so, where the latter category encompasses any piece of fake news that is accidentally true.[10] This argument by analogy is a non-starter. On the standard definition of lying, in fact, 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 and 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. If lying tolerates the assertion of a true proposition, the analogy between telling a lie and spreading fake news, far from demonstrating that fake news must lack truth, encourages the opposite thought that fake news can be true, and therefore makes room for 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 relationship between the available evidence and the truth of the relevant 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 conceded that they would 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.

---

[9] See Mukerji (2018:) for an opposite analysis of the case.
[10] Jaster and Lanius's parallel also considers the case of murdering, which we omit for the sake of brevity.

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.[11] More specifically, 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 Jaster and Lanius's sense, but that its creator takes it to lack sufficient evidence in favor of the truth of their assertion.[12]

The analysis of the properties of fake news that we have offered in this section motivates the following account of fake news:

> By asserting/implicating that P, a subject S is asserting/implicating a piece of fake news if and only if (i) S's assertion/implication that P is meant to address a large enough audience[13] (ii) in the guise of news; (iii) P is asserted/implicated either with the intention to deceive one's audience believe or with no concern with P's truth; and (iv) S does not think that P is supported by the evidence.

It might be objected that our account is overly restrictive, in that it does not take into consideration that also the attitude with which a content is re-asserted in the process of its propagation, or the evidence based on which it is reasserted, can matter to its status. This complaint can be substantiated in light of examples adduced by Pepp, Michaelson, and Sterken (2019b: 88-9) to illustrate an alleged virtue of their *dynamic* account, namely the idea that what at a particular time is (not) fake news can become (cease to be) fake news at a later time. In one of these examples, a piece of fake news is originally asserted with the right intentional and evidential properties—say with the intention to deceive and based on no evidence—but it intuitively ceases to count as fake

---

[11] It is worth emphasizing that Jaster and Lanius (2018: 220) 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 and Lanius commit themselves to denying that they amount to pieces of fake news, our analysis explains why it is reasonable to consider cases of this sort as genuine fake news.

[12] We take it that E is sufficient evidence in favor of P just if E is a comprehensive body of evidence on the basis of which S could justifiably believe that P. It is important to note that S can fail to have sufficient evidence to believe that P and believe, nonetheless, that she has sufficient evidence to believe that P. This can happen, for instance, when S believes she has sufficient evidence in favor of P, in fact S's evidence supports P, but this evidence is incomplete. In a similar case, S's assertion of P can be criticised as careless or too hasty, but not as the assertion of a piece of fake news.

[13] One might worry that the concept of a 'large enough' audience is too vague. In response, we shall point out that the vagueness of the definition of fake news parallels the vagueness of the definition of news in general. In a word, we take a large enough audience to be the audience of an assertion we would typically consider to be a piece of news.

news when it is taken up by a diligent journalist who investigates the matter and, having the misfortune of encountering bad sources which confirm it, re-asserts the news with the intention to be reporting a fact. A different example, also inspired by Pepp and colleagues, illustrates the same complaint, but in the opposite direction.[14] In this example, a purely satirical piece is reposted with the more malevolent intent to deceive the audience in a context in which its original spirit is no longer discernible. In this case too, it might appear natural to say that a change in status has occurred and that something that was not fake news has now become fake news.

Far from representing a threat to our account, both cases are easily accommodated by it. First of all, a natural way of describing these examples is precisely to say that a piece of fake news (or of something that was not fake news) has been *mistaken* or *misused* by its propagator as something that was not (or was) fake news. This way of looking at the examples confirms our account, as it lets the real status of the content depend on the intentional and evidential properties of its original assertor. The honest journalist of the first example can be described precisely in these terms, as committing the unfortunate mistake of divulging a piece of fake news as real news. In the same spirit, it seems appropriate to regard the malevolent propagator in the second example as fraudulently appropriating satirical content and misusing it as if it were fake news.

Granted, there is also a way of looking at these examples according to which what the propagator has asserted has become something different in the process, as Pepp and colleagues contend. On closer inspection, however, also this way of describing the examples is compatible with our account. To look at the article as a piece of news, or at the satirical piece as having become fake news, in fact, is to succumb to the temptation to regard the honest journalist or the malevolent propagandist as *initiating* entirely new processes of propagation.[15] It is because we treat—with some reason, given the difference in intentional and epistemic properties—them as new originators, and so we let *their* intentional and epistemic properties matter to the content's status, that we are no longer inclined to look at the contents involved in the same way. So, again we are resting our reading of the example on our account.

---

[14] The original example takes its inspiration from an historical fact. An *Onion* satirical piece said that rural Americans would have preferred to vote for Ahmadinejad rather than voting for Obama. Reportedly, this piece has been *mistaken* by an Iranian news agency and divulged as real news.

[15] This temptation would be admittedly even stronger in the case in which a content that on our account was an *accidentally true* piece of fake news started to circulate accompanied by newly discovered evidence of its truth. In this case, we would say that a new process of propagation has begun.

3.   On the intention to mislead one's audience: The case of clickbaiters

As we have seen, a crucial aspect of our account is that the intention to deceive one's audience is not a necessary condition for a content to be a piece of fake news. We agree with Jaster and Lanius (2018), Mukerji (2018), and Gelfert (2018) that fake news can also be asserted with *no concern* for the truth of what one asserts. In this case, the assertor might be deceiving one's audience, but need not be pursuing this goal, or so we shall argue in the remainder of this section. Pritchard (*forthcoming*) has recently argued against this idea in his discussion of Gelfert's view.[16] Since his objection calls into question the third condition of the account we have just proposed, we shall attempt to show how the objection can be met. The discussion of this aspect of the notion of fake news will reveal a related—in fact, more fundamental—problem of the current debate in the epistemology of fake news. We all agree that consuming fake news is detrimental to our epistemic wellbeing. However, the main epistemic threat of fake news consumption is not that it makes us believe false or misleading claims, but rather that it fosters our disinterestedness in the truth and makes us prone to consume whatever news satisfies our partisan attitudes and need for social recognition.

Let us go back to the case of the Macedonian clickbait factories and consider it more closely. It is widely agreed that the contents fabricated and disseminated by clickbaiters are highly misleading and often amount to fake news; the controversy lies in the diagnoses that have been offered as to why this is the case. As anticipated in the previous section, if they are merely after making profits out of the traffic they are able to generate with their contents, it is not clear that their attitude should include an intention to deceive their audience. Rather, it might well just be that they have no concern for the truth of what they share.

According to Pritchard, this diagnosis of the case neglects the obvious possibility that the creation and on-line distribution of misleading contents may be prompted by more than one motive. Pritchard concedes that the overarching motivation behind the clickbait industry is the generation of financial revenues. However, he also contends that this goal requires the pursuit of an intermediate and instrumental goal, that of persuading prospective readers of the likely truth of the contents on offer. Otherwise, asks Pritchard rhetorically, "why on earth would they be clicking

---

[16] Pritchard's second criticism, which we do endorse, pertains to Gelfert's characterization of fake news as a degenerate species of the genus news. Pritchard observes, correctly in our opinion, that Gelfert objectionably derives a metaphysical claim pertaining to what fake news is—a type of news, according to him—from an epistemological claim––pertaining to the difficulty of telling genuine news from fake news. That Gelfert should find this transition compelling is especially surprising in the context of his paper. Drawing a parallel between the couple information/disinformation and the couple news/fake news Gelfert in fact quotes Floridi (2016), who, in fact, is explicit that false news is no news at all.

on them under the guise of them being news?" (*forthcoming*: 4).[17] Thus, on Pritchard's view, clickbait farmers must be pursuing the additional goal to deceive their readership.

We disagree with Pritchard's diagnosis of the clickbait factories case: the accomplishment of clickbaiters' financial goal relies on their ability to induce a twofold reaction on the part of consumers, namely that people click on their contents and repost them widely. At no point in this chain, clickbaiters "must be in the business of also wanting to actually mislead" (*ibid.*) their clientele, or so we shall argue.

Before moving on, however, it is necessary to shed light on a potential terminological ambiguity concerning what it means to 'mislead' or 'deceive' one's audience. As seen above, we deny that fake news producers invariably pursue the goal of making their audience believe false or misleading contents. In this *strong sense* of deceiving, we deny that being asserted with the intention to deceive is a necessary condition of fake news. We are ready to admit, though, that in a *weak sense* fake news is always asserted with the intention to deceive. One weaker way in which fake news producers want to deceive their audience has to do with their attitude, in that they must at least *pretend* to be talking seriously (Mukerji 2018). Another way in which fake news producers deceive their audience in a weak sense has to do with the presentation of their contents. Surely, the stories they circulate, however preposterous or outlandish, must be crafted to look admissible candidates for being real news reports: incredible or explicitly fictional stories would not suit the purpose of generating ad revenues. Thus, if Pritchard's contention is simply that fake news must be deceitful in any (or both) of these weaker senses, our disagreement is largely a matter of terminology.

Yet, if Pritchard's view is that the intention of fake news producers is one of deceiving their audience in the strong sense, then it looks as though the clickbait case possibly speaks against his view. The general picture presupposed by Pritchard seems to be that fake news producers target an audience hungry for reliable information and aim to meet their demand by selling a fraudulent counterfeit of the desired epistemic good. As we shall see in a moment, however, a more appropriate story has it that the consumption and sharing of fake news are underwritten by the pursuit of such aims as strengthening one's political identity, tightening social bonds with partisan fellows, and increasing one's sense of belongingness to a wide (sort of online) community, which do not necessarily involve believing the contents one contributes to spread. More straightforwardly, it looks as though clicking and reposting fake news can be explained also by a different sort of hunger, namely a hunger for social recognition.

---

[17] Rini makes a similar point, when she observes that "people who make money from fake news are perfectly happy if nine-in-ten of their readers are not deceived, but they *do* need some percentage to be deceived long enough to convey the link to future clickers" (2017, p. E-45).

If this is right, we think that a conceivable (and perhaps even true) story can be told in which clickbaiters, aware of the mechanisms underwriting the consumption of fake news online, pursue their financial goal without designing their fakes to be deceitful in the strong sense (namely, with the intention to induce false beliefs). In the (possible or actual) scenario we have in mind, fake news producers realize that carefully fabricating stories that will eventually be able to persuade their readers would be unnecessarily time-consuming and effortful (and so economically inefficient). So they publish their stories without believing that they will convince their readers. They just make sure that, apart from being minimally plausible (namely deceitful in a weak sense), these stories satisfy non-epistemic desiderata (basically, being the kind of stories one and one's partisan fellows would like to be true) with the intention, by so doing, to bring their potential readers to consume and share them independently of whether the readers are committed to their truth. We think that in a similar scenario the alleged piece of information being manufactured and shared online, provided it is not explicitly presented as fiction, would still be fake news.

The hypothesis that grounds our conjecture—that fake news is consumed as a way of satisfying a hunger for social recognition—is supported by empirical evidence from psychology and business studies. As Berger and Milkman (2012) have found out, virality is (at least partly) driven by emotional response and, in particular, it looks as though we mostly share contents we feel strongly about. As a recent briefing on the psychology of online sharing of misinformation shows (Vicol 2020), this feeling is grounded in the dynamics underlying the construction of our social identities as members of online communities: "feeling strongly about a topic is not just an individual reaction to a strong emotion. It is also a social process of affirming what we stand for, who we like and what we wish to be perceived as" (2020: 14).

Furthermore, in a study on the associations of the dark side of social media use and fake news sharing behavior, Taiwar and colleagues have recently found that one's proneness and likelihood to share fake news can be increased by several factors that trigger social and psychological aspects of our social media agency rather than epistemic ones. For example, the *fear of missing out*—i.e., the anxiety that other people are living experiences from which one is excluded and that therefore generates the need of being in constant contact with them (Przybylski et al. 2013)—makes people willing to repost fake news, insofar as one thinks that one's contacts will see, and interact with, the content one shares. Similarly, Taiwar and colleagues argue that proneness to share fake news correlates positively with *high self-disclosure*, namely with one's willingness to divulge personal information with others. Those subjects will be more likely to share fake news because sharing contents with no regard to their epistemic pedigree can fulfil the same need "to increase mutual understanding, strengthening relationships, and enhancing bonds among group members" (2019:

75) that characterizes self-disclosed individuals.[18] In sum: these factors indicate that people might be willing to trade in sharing fake news as a way to build up their social identity, that is, to establish or reinforce their social bonds with other users and feel part of an online community. None of this requires that people believe the fake news they consume and share. Rather, by shedding light on our proneness to cultivating the aforementioned non-epistemic ends, this analysis reveals that the main epistemic threat of fake news consumption lies in the fact that it distracts us from the truth (and the other epistemic features) of the contents we engage with.

Granted, acknowledging that these non-epistemic factors contribute to making one willing to share fake news on-line is not yet to deny that this sort of behavior may involve genuine *beliefs*. It is surely the case that often fake news consumers are after reinforcing both their beliefs and their social identities.[19] However, all that matters for our purposes is to show that, in a relevant amount of cases, people consume and share fake news as a way to satisfy the latter goal, regardless of the former.

This point may be brought into sharper relief by pointing to possible cases in which, while the recipients of a specific piece of fake news do not believe it, they nonetheless accomplish the practical needs reviewed above by consuming and sharing it widely. This happens whenever simply accepting a piece of fake news and communicating it online proves instrumental to the fulfilment of, say, the need to reinforce the social bonds with one's fellow partisans or increase one's feeling to belong to a wider community. As a label for this specific attitude we propose "treating a content as settled for the purpose of social recognition". When we say that one displays this attitude toward a content, we do not mean that one *poses* as though one believes a content one has already decided not to believe. Nor do we mean that one cynically deploys premises believed to be false or epistemically problematic for the sole purpose of achieving some dialectical advantage in argumentation.

What we have in mind is rather an attitude that precedes deliberation, is largely unconstrained by the evidence, and is mainly responsive to practical considerations. We take this attitude to be mostly *unreflective*, that is, a spontaneous and intellectually less sophisticated response to the encounter with contents one perceives to be fitting with one's most cherished convictions. It does not seem hard to imagine a case in which a supporter of a political party accepts a piece of fake news reporting the alleged misconduct of a politician of the opposite party with this attitude. In a

---

[18] See also Whitty and Joinson (2009).

[19] Well-known psychological mechanisms such as the confirmation bias or motivated reasoning may, in fact, dispose one to form the beliefs that best suit the need of a strengthened political identity, or enhanced social integration, either by bringing one to attend selectively to the evidence that supports them or by inclining one to cast the available evidence in a more favourable light (Lazer et al 2018).

case like this, the character of our example does not bother checking the reliability of the information simply because they are moved by the implicit appreciation that it is the kind of thing that, when sufficiently shared, would tend to reaffirm the sense of the superiority of their political affiliation, thereby reinforcing it. In this case, they will possibly treat the content as true, and repost it on-line, as a way of strengthening their identity qua decent member of their community and independently of a real commitment to its truth.

4. *Concluding Remarks: A Pathological Relationship with Information*

In this chapter, we have defended a novel account of what fake news is. We take it that our account has offered a twofold contribution to the debate on the epistemology of fake news. For one thing, it has warned us not to postulate too intimate a relation between the spread of fake news and the intention to deceive. We have suggested that the process through which fake news is distributed should not be merely regarded as an exchange of epistemic goods, but rather as a way for social media users to build up their social and political identity. As we put it, fake news consumption often satisfies one's hunger for *social recognition* more than reliable information. A consequence of our suggestion is that a different attitude than belief—which we have labelled "treating a content as settled for the purpose of social recognition"—might often be implicated within the practice of consuming and sharing fake news online.

At a more general level, the discussion offered in this chapter sheds light on a second and more profound threat that has largely escaped the recent attention paid to fake news. If we are right on the reason why fake news producers need not aim to deceive (in the strong sense) their audience, fake news turns out to be epistemically problematic not only for the consequences it engenders—the spread of false or misleading beliefs—but also, at a more basic level, as a symptom of the pathological relationship with information that underwrites its proliferation and that cannot but be exacerbated by its uncontrolled diffusion. Thus, if what we have argued about the attitudes of fake news consumers is right, there seems to be scope for retargeting the discussion about the epistemic problems generated by the spread of misinformation. For the crux of the matter is less about the fact that social media users are not good at selecting products that can satisfy their hunger for reliable information and more about the fact that they lack the appropriate epistemic attitudes toward information, due to their hunger for social recognition. This is particularly relevant if we care about counteracting the spread of misinformation. For our analysis suggests that a widely endorsed reparative approach, aimed at providing consumers with more information and teaching them how to spot misinformation, can hardly prove to be the right therapy.

References

Allcott, H. and Gentzkow, M. (2017). Social Media and Fake News in the 2016 Election, *National Bureau of Economic Research*, Working Paper 23089.

Berger, J., Milkman, K. L. (2012). What Makes Online Content Viral? *Journal oof Marketing Research* 49(2): 192-205.

Chisholm, R. M., and Feehan, T. D. (1977). The Intent to Deceive. *Journal of Philosophy* 74: 143–159.

Croce, M., Piazza, T. (2020). Epistemic Vices and Responsibility in Fake News Consumption. Unpublished manuscript.

Dentith, M. R. X. (2017). The Problem of Fake News. *Public Reason* 8(1-2): 65–79.

Floridi, L. (2011). *The Philosophy of Information.* Oxford: Oxford University Press.

Frankfurt, H. G. (2005). *On Bullshit.* Princeton: Princeton University Press.

Gelfert, A. (2018). Fake News: A Definition. *Informal Logic* 38(1): 84–117.

Gettier, E. (1963). Is Justified True Belief Knowledge?. *Analysis* 23(6): 121–123.

Habgood-Coote, J. (2018). Stop Talking About Fake News. *Inquiry 62 (9-10)*: 1033-1065.

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., Lanius, D., (2018). What is Fake News?. *Versus* 2(127): 207–227.

Klein, D., 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. J., Baum, M. A., Benkier, Y., Berinski, A. J., Greenhill, K. M., Menczer, F., … Schudson, M. (2018). The Science of Fake News. *Science* 359(6380): 1094-1096.

Mahon, J. E. (2008). Two Definitions of Lying. *International Journal of Applied Philosophy* 22: 211–230.

Mukerji, N. S. (2018). What is Fake News?. *Ergo: An Open Access Journal of Philosophy* 5(35): 923-946.

Pepp, J., Michaelson, E., Sterken, R. K. (2019a). Why We Should Keep Talking about Fake News. *Inquiry*. Doi: 10.1080/0020174X-2019.1685231.

Pepp, J., Michaelson, E., Sterken, R. K. (2019b). What's News About Fake News? *Journal of Ethics and Social Philosophy* 16(2): 67-94.

Piazza, T., Croce, M. (2019). Epistemologia delle fake news. *Sistemi Intelligenti* 31(3): 33-461.

Primoratz, I. (1984). Lying and the "Methods of Ethics". *International Studies in Philosophy, 16*, pp. 35–57.

Pritchard, D. *Forthcoming.* Good News, Bad News, Fake News. In S. Bernecker, A. Floweree, T. Grundmann (eds.), *The Epistemology of Fake News*, Oxford: Oxford University Press.

Przybylski, A. K., Murayama, K., DeHaan, C. R., Gladwell, V. (2013). Motivational, Emotional, and Behavioral Correlates of Fear of Missing Out. *Computers in Human Behavior* 29(4): 1841-1848.

Rini, R. (2017). Fake News and Partisan Epistemology. *Kennedy Institute of Ethics Journal* 27(S2): E43–E64.

Taiwar, S., Dhir, A., Kaur, P., Zafar, N., Alrasheedy, M. (2019). Why Do People Share Fake News? Associations between The Dark Side of Social Media Use and Fake News Sharing Behavior. *Journal of Retailing and Consumer Services* 51: 72-82.

Vicol, D. O. (2020). Who Is Most Likely to Believe and to Share Misinformation? *Africa Check, Chequeado, and Full Fact.*

Whitty, M., Joinson, A. (2009). *Truth, Lies, and Trust on the Internet.* London: Routledge.

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