

CONSUMING FAKE NEWS: CAN WE DO ANY BETTER?

MICHEL CROCE

University College Dublin

TOMMASO PIAZZA

University of Pavia

ABSTRACT. This paper focuses on extant approaches to counteract the consumption of fake news online. Proponents of structural approaches suggest that our proneness to consuming fake news could only be reduced by reshaping the architecture of online environments. Proponents of educational approaches suggest that fake news consumers should be empowered to improve their epistemic agency. In this paper, we address a question that is relevant to this debate: namely, whether fake news consumers commit mistakes for which they can be criticized and that they could easily avoid by reforming their doxastic conduct. Proponents of structural approaches, like R. Rini and B. Millar, have defended in different ways a negative answer to this question. In this paper, we criticize their views and suggest that individual users could improve on their epistemic practice by widening their information diet.

KEYWORDS: fake news; epistemic partisanship; cognitive biases; epistemic blame

[preprint version]

1. INTRODUCTION: FAKE NEWS CONSUMPTION AND AVAILABLE REMEDIES

Social media such as Facebook and Twitter are virtual environments in which misinformation abounds and, owing to the activities of their members, circulates widely and is often met with widespread acceptance (Vosoughi et al. 2018). The complex dynamics that involve sharing and believing fake news, which can be epitomized as *fake news consumption*, make it harder for individual agents to be accurately informed and, in turn, negatively affect the well-being of our communities at a social and political level. Thus, fake news consumption is almost unanimously regarded as a disease for which we need urgent therapy (Lazer et al. 2018), and many scholars from multiple disciplines have started looking for suitable remedies.

This paper focuses on extant approaches to counteract the consumption of fake news. Before introducing them, it is important to acknowledge that the very notion of fake news is the locus of a hot discussion in the philosophical literature (e.g., Dentith 2017; Gelfert 2018; Habgood-Coote 2018; Jaster & Lanius 2018; Mukerji 2018; Pepp, Michaelson, and Sterken 2019; Pritchard 2021). We have proposed our own favored account of fake news elsewhere (Croce & Piazza 2021), but for the purposes of this paper, we will leave the notion of *fake news* at the intuitive level, as illustrated by standard cases featuring false or misleading contents presented as news and not conveyed with the intention to inform one's audience. (Think, as an example, of the preposterous news reports circulated right before the 2016 US Presidential elections —often referred to as the Pizzagate—claiming that Hillary Clinton was involved in a child trafficking ring.)

At present, proposed remedies to the consumption of fake news tend to distribute in two distinct categories. According to *structural approaches*, it is the very nature of the online environments we inhabit that brings so many people to believe and share fake news. Proponents of these approaches suggest that the architecture of online environments should be reshaped to reduce our proneness to consuming fake news. One possibility is to implement a system that signals uncertified sources of information as such (Lazer et al. 2018; Vosoughi et al. 2018). A different proposal is to disincentivize uncritical sharing by letting the testimonial reputation of individual users vary inversely with the quantity of misinformation one has contributed to circulate (Rini 2017: E-57).

Other theorists put more emphasis on reforming doxastic conduct at the individual level by suggesting that fake news consumers should be empowered to improve their epistemic agency. Strengthening citizens' internet literacy skills (Lewandowsky et al. 2012) is an example of a similar proposal. Fostering their intellectual virtues (Heersmink 2018; Smart 2018) is another. As these approaches focus on individuals and their doxastic conduct rather than reshaping the environment they inhabit, we shall call them *educational approaches*.

Proponents of systemic approaches often remark that fake news consumers—i.e., the real target of educational approaches—do not commit mistakes for which they could be appropriately criticized.

Granted, technically speaking, systemic and educational approaches are compatible with one another. A mix of both approaches could offer a promising recipe against the spread of fake news, and some theorists may well be propounding hybrid remedies of exactly this kind. However, the emphasis that many proponents of systemic approaches place on a favorable assessment of the doxastic conduct of fake news consumers might easily serve the opposite goal of *undercutting* the motivation of educational approaches. This can be noted by inspecting the following argument.

- (1) Since educational approaches aim to correct and ameliorate the epistemic conduct of social media users, they presuppose that social media users commit mistakes that can be amended.
- (2) This diagnosis is wrong: While it is not a desirable thing that social media users believe fake news, for the most part, it is not a result of epistemically criticisable conduct either.
- (3) Thus, educational approaches are based on poor motivation: The only way to promote the epistemic well-being of ordinary consumers of fake news is to intervene in the structural features of the information environments they inhabit.

In this paper, we shall not take issue with the first premise: presumably, many would be willing to concede that educational approaches presuppose that social media users commit correctable mistakes. Given the initial plausibility of the first premise, any argument supporting the second premise appears to strengthen the argument's conclusion.

We will focus, in particular, on two possible ways in which the second premise of the above argument has been motivated. The first defense is provided by R. Rini (2017), who has suggested that fake news consumption, rather than by criticisable doxastic routines, may be underwritten by the exercise of an intellectual *virtue* that she calls "epistemic partisanship." The second is supplied by B. Millar (2019), who has defended the view that social media users are fully excused by their cognitive biases and the way in which social media algorithms amplify the effects of our cognitive limitations.

It is unclear whether Rini and Millar endorse the argument above because they have not explicitly engaged with the first premise. So, it may well be the case that they would consent to a hybrid approach based on both systemic and educational elements. However, it also deserves emphasis that, above and beyond defending the second premise, they have come close to explicitly endorsing its conclusion. For Rini, finding a suitable solution to the problem of fake news requires that we “look beyond individual epistemic practices—to institutions” (2017: E-54). For Millar, “the only way to alleviate the current epidemic of false beliefs is to change the information environment itself” (534), and instead of focusing on individuals, “we should focus on the technologies and public policies that have created an information environment that makes widespread false beliefs inevitable” (Millar 2019: 526).

Be that as it may, providing a comprehensive interpretation of these views is not our primary concern in this paper. Our aim is rather to shield the claim that educational approaches can offer legitimate help in contrasting fake news from the argument above. For this reason, in the remainder of this paper, we shall address Rini’s and Millar’s defenses of the second premise of the argument with the sole goal of showing how its conclusion could be resisted. We shall leave it open whether Rini and Millar actually subscribe to the argument and endorse its unfriendly conclusion to educational approaches.

The paper is structured as follows. Section 2 attempts to undermine Rini’s thesis that the consumption of fake news often arises out of praiseworthy behavior on the part of social media users. Section 3 critically addresses Millar’s thesis that, if not epistemically virtuous, social media users can nonetheless be excused for believing fake news by the combined effect of their proneness to cognitive biases and the structural features of the online environments they inhabit. In section 4, we draw our conclusions.

2. FAKE NEWS AND EPISTEMIC PARTISANSHIP

Educational approaches seem to presuppose that epistemic agents consuming fake news do so because they commit mistakes for which they can be criticized and which they can be asked (or can be educated) to amend. This section takes issue with a recent view put forth by Rini (2017), according to which most fake news consumption is sustained by partisan attitudes. For this reason, not only is it the case that social media users do not commit any recognizable mistakes for which they can be criticized, but also, their doxastic conduct can be praised as epistemically virtuous.

In Rini's view, epistemic partisanship is the intellectual virtue one exercises when assigning greater credibility to the testimony of one's friends—or one's "co-partisan fellows"—in normative domains (E-50). More specifically, Rini has in mind two classes of claims. The first class encompasses political or moral claims, which directly voice one's commitment to fundamental values, such as freedom, equality, autonomy, etc. An example of a normative claim is the claim that abortion is morally wrong. The second class encompasses what Rini calls *politically* and *morally relevant* claims, which amount to descriptive claims relevant to political or moral evaluation. As an example of a politically relevant claim, Rini mentions the alleged report about the misconduct of a political candidate, which is relevant to the assessment of her character.

Rini's claim that epistemic partisanship is an intellectual virtue is relevant to the normative assessment of the epistemic conduct of fake news consumers because the communities to which we belong—like the network of our friends and contacts on Facebook or Twitter— are often *partisan* (Van Bavel & Pereira 2018). One obvious consequence of the structure of these online communities is that most beliefs we acquire there—including beliefs in fake news—are formed based on the testimony of like-minded individuals—i.e., our co-partisan fellows. Thus, if Rini is right these are beliefs we have formed by exercising the intellectual virtue of epistemic partisanship, and we cannot be criticized for having them. So, the crucial question is whether Rini is right in describing epistemic partisanship as an intellectual virtue in the relevant range of cases.

The problem with Rini's argument lies with how she conceives of the sort of claims that typically spread fake news. We can agree with her that insofar as we have reasons to believe that our contacts share our value commitments, *all other things being equal*, we are justified in trusting their normative claims (or other parties' normative claims, if our contacts agree with them). However, fake news rarely pertains to normative claims in the above sense. It more likely pertains to normatively relevant claims, namely, descriptive claims that have a clear bearing on some normative questions.

Why should we follow Rini and believe that we would be rational in assigning greater credibility to the reports of our co-partisans also when they pertain to descriptive domains? In her view, partisan affiliation can be relied upon also in assessing "testimony about descriptive claims," at least when these are politically related, "because the act of transmitting political news implicates *normative decisions* on the part of the testifier. Often these are decisions about what is politically *important*" (2017: E-52). Let's try and unpack the suggestion on offer.

On one, possibly uncharitable reading, Rini might be taken to suggest that when our co-partisans share a normatively relevant claim P we are licensed *not only* to expect P to be *important* but also, to the extent to which we take it to be important, to think that P is likely true. In this reading, Rini's suggestion can be discarded right from the start. The question we are faced with when we receive testimony that P is not whether we are rational in deciding to spend time consuming the relevant news report. Rather, the question is 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 is no reason to suppose that they sustain an affirmative answer to the second as well.

However, Rini might be read more charitably. To introduce this alternative reading, we have to keep the normatively relevant claim P, reported by our co-partisan fellow, distinct from the normative claim Q, for which P is normatively relevant. Suppose P is the claim that a political candidate didn't pay her taxes twenty years ago. While P is a descriptive claim, it is clear that it has some bearing on the normative proposition (Q) that the political candidate is morally corrupt. In this sense, Q is a proposition for which P is normatively relevant. Of course, it may be an open question to what extent

one should increase one's credence in Q, given that one has learned P. Perhaps Rini is merely saying that the degree to which one should increase one's credence in Q given that one has learned P is greater when P has been reported by a co-partisan fellow than it would have been otherwise. For instance, she might be suggesting that when a co-partisan fellow shares the information that the candidate didn't pay her taxes, since I am entitled to think that my fellow has a good sense of what is normatively relevant, I am also thereby entitled to take this fact not to be a minor lapse on the part of the candidate but to be highly revealing of her moral character.

If this is true, Rini might be read as putting forward the following general maxim. When P is a normatively relevant claim, and Q is a normative claim for which P is normatively relevant, the degree to which one should increase one's confidence in Q given that one has learned P varies with the normative distance between one and the source of the testimony that P: the closer the testifier, the more confident that Q one should become after learning P.

This second reading of Rini's suggestion is possibly more accurate than the first. However, we believe that it proves equally unhelpful to vindicate *as intellectually virtuous* the practice of believing in fake news when we receive it from our co-partisan network, as it often happens in social media. When my co-partisan peer reports that a political candidate didn't pay her taxes twenty years ago, the awareness that it was my co-partisan fellow who shared this report may well explain why I am epistemically entitled to credit the alleged fact *itself*—the political candidate's tax fraud—with a greater significance concerning the question of what kind of person she really is. However, it does not explain why I should also be epistemically entitled to overinflate my confidence that she didn't pay her taxes. So, it fails to show that the practice under scrutiny—namely, the practice of accepting a piece of fake news based on the testimony of one's co-partisan peers—is epistemically virtuous¹.

¹ As an anonymous referee of this journal has pointed out to us, it is sometimes argued that it is reasonable to place greater trust in the testimony of individuals who belong to groups with which one identifies. The reason would be that, since these individuals are likely to share our interests, they are less likely to try to deceive us. (For similar arguments see, e.g., Levy 2019: 321-322). According to this referee, arguments of this sort could offer another route to vindicating the intellectual virtue of epistemic partisanship when it comes to politically related testimony. Although we find it an interesting suggestion, we also believe that it would ultimately fail as a defense of Rini's central claim. For the mere fact that people sharing our same interests are less likely to try to deceive us does not imply that they also are a reliable route to the truth. It is perfectly compatible with them also being unreliable due to some epistemic shortcoming.

If our analysis is correct, Rini's argument fails to show social media users believing pieces of fake news received from their co-partisans instantiate an epistemically virtuous behavior. Quite to the contrary, it is a moral of our rejoinder to Rini's view that their conduct is epistemically criticisable because it reveals their inability to distinguish the factual aspects from the normative aspects in politically relevant claims and, in turn, their inability to accept these claims as true based on the former considerations, rather than the latter. Any educational approach aiming to help social media users assess the credibility of information sources should strengthen their ability to discern the normative and descriptive components of news reports as well as the pros and cons of interacting with proximate and like-minded sources.

3. COGNITIVE BIASES AND EPISTEMIC BLAMEWORTHINESS

By rejecting Rini's view that fake news consumers might be intellectually virtuous, the previous section has blocked a potentially worrisome version of the argument according to which educational approaches are unnecessary. In this section, we shall complete the job by taking issue with Millar's view of fake news consumers as blameless epistemic agents. Although the main target of Millar's project is a normative assessment of the epistemic conduct of social media users, his view is key to the goals of this paper. For, if Millar is right that social media users are fully excused by the relevant features of their environment and by how these features exacerbate their cognitive limitations, the second premise of the argument introduced in §1 would be successfully vindicated, and its conclusion made more difficult to resist. For this reason, in the remainder of this section, we shall try and put pressure on his assessment of the epistemic performances of social media users.

For starters, Millar argues that fake news consumers are blameless because their proneness to misinformation can be typically explained by human proneness to cognitive biases and by the way in

which the structure of social media amplifies the effects of these biases. Millar focuses on two specific biases, namely the truth effect and confirmation bias.

The *truth effect* amounts to our tendency to form the belief that P when P is repeated many times (Fazio et al. 2015; Levy 2017). This tendency is key to explaining fake news consumption on social media. In these environments, as we have illustrated in the previous section, users typically cluster around networks of like-minded people. For this reason, the same pieces of misinformation will likely appear multiple times within one's newsfeed, due to repeated sharing by members of the network. This will highly enhance the probability that social media users, given the general tendency to accept what is repeated many times, end up believing it.

Empirical evidence (Nickerson 1998) also shows that we are hard-wired to form the belief that P less critically when P coheres well with our pre-existing beliefs. This is known as *confirmation bias*. Confirmation bias plays an equally decisive role in explaining how social media users get to believe in fake news because the architecture of social media takes advantage of this tendency to make us spend our time online. More precisely, social media are run by algorithms—typically opaque and undetectable by users—designed to ensure that we are presented with news reports that align with our views on political, social, and moral matters (Bozdag 2013; Nichols 2017; Pariser 2011).

These considerations about the limitations of our cognitive capacities—and the ways in which social media exploit them—suffice to explain why it is especially difficult, from a psychological point of view, to refrain from believing the pieces of fake news to which we are exposed in online environments. More precisely, in Millar's view, these factors have such a negative influence on our capacity to assess the quality of the information we consume that it would be unreasonable and unrealistic to require that social media users refrain from believing in the fake news they encounter.

Although the influence of the psychological and environmental factors explains why social media users cannot but form beliefs in fake news, it might well be the case that they can exert some control over these factors. If this were true, social media users could be criticized, if not for their beliefs in fake news *per se*, then at least for their failure to control the forces that make those beliefs

psychologically unavoidable. Millar admits that “our tendency to hold onto our beliefs is not so strict that we are incapable of changing our minds” (2019: 529), but he also argues that keeping our psychological tendencies under control requires so much effort that we are excused if we fail to do so—viz., an *excuse of force*, in his terms.

In particular, according to Millar the most promising ways to mitigate—if not overturn—the negative effects of our cognitive biases would require exposing us to counterevidence for the beliefs we form under the influence of these biases, or, more directly, putting ourselves out of range of their influence by avoiding social media altogether. The problem with the first suggestion, however, is that collecting a sufficient amount of counterevidence for Millar requires that one actively fact-checks “potentially false claims by seeking out additional information from alternative sources” (530), and it would again seem unreasonable to expect that social media users have the time, energy, and resources to undertake this process for any candidate belief they are fed within these environments. As regards the second suggestion, Millar observes that being a member of some social network or another has become an unavoidable feature of our social life, and, for this reason, concludes that it would be unrealistic to advance the request that social media users counteract their proneness to consume fake news by simply *ceasing to be* social media users.

If Millar is right, then it looks as though social media users cannot be blamed for how they manage their online epistemic conduct. We cannot expect that they avoid forming beliefs in fake news because cognitive biases make it very hard for them to distinguish good news from bad news. But more importantly, we cannot even ask that they do better at distinguishing good news from bad news by seeking out additional evidence, or by avoiding problematic sources, because they lack the resources to get the job done consistently.

We tend to agree with Millar on both counts. However, we also believe that a different strategy to ensure that social media users get better at distinguishing good from bad news is available and deserves to be implemented. In short, this is the strategy of having social media users widen their sources of information, at least by consuming news from traditional media such as newspapers and

a sensationalist piece reporting Pope Francis' endorsement of Donald Trump before the 2016 US election. Suppose you know that this piece of news has not been previously reported by the mainstream media, as you are accustomed to a varied information diet. Unless you have reasons for thinking that this omission resulted from negligence or the news source's investigative limitations, it provides you with an indication that the piece is fake; that is, with a reason to not accept the piece as real news. In general, the fact that the mainstream media didn't report a piece of alleged news that, if true, they would have published, provides indirect evidence that the piece is fake news. For it indicates that the relevant media have evidence that it is fake, or at least that they were not able to find evidence that it was true news.

However, several concerns can be raised against our suggestion.⁵ Firstly, one might worry that the mainstream media contribute themselves to spreading fake news. This might happen indirectly, when they report the words of people who are themselves voicing fake news; or directly, as in the case of partisan media (such as Breitbart or Fox News) that smuggle in misinformation as real news. If this is true, the recipe of consuming mainstream media may result in the acquisition of more fake news or in reinforcing belief in fake news encountered elsewhere. Secondly, one might worry that the remedy we are recommending will prove useless in all cases in which people will interpret the omission of some alleged piece of news by the mainstream media in a prejudicial way, as the result of a malevolent intention to deceive one's audience or conceal some important fact which does not fit their political agenda. Thirdly, one might contend that the social media users that constitute our target already consume mainstream media. In this scenario, the plausibility of our account would hinge on the availability of evidence that a more varied information diet would diminish the proliferation of fake news. Since we are unable to point to the right kind of evidence, the objection goes, our proposal would be ultimately unsupported.

that the mainstream media will often remedy their initial omission by starting to cover the relevant news. In sum, we expect that the cases in which by adopting Goldberg's recipe one will dodge fake news greatly outnumber the cases in which, by so doing, one will end up missing out on true news.

⁵ We owe a detailed presentation of these possible worries to an anonymous referee of this journal.

While these objections would deserve a lengthier discussion than we can provide in this paper, the following considerations motivate the optimistic expectation that they can be satisfactorily dealt with.

The first objection sheds light on the difficult problem for mainstream media of finding a balance between the duty to report what relevant public figures maintain—including cases in which what they assert is mostly fake news—and the duty to inform their audience, that is, to provide them with high-quality information⁶. This problem becomes even bigger if media outlets themselves are involved in distributing misinformation. As anticipated, this typically happens with partisan media such as Breitbart and Fox News.

The underlying assumption of this objection is that these partisan media belong within the same broad category of mainstream media as more balanced news outlets. If this were true, enlarging the information diet in the way we suggest would put social media users at a high risk of being poisoned, informationally speaking, precisely by those sources that are supposed to protect them from the spread of misinformation. But we see no reason to grant this presupposition. On the contrary, we agree with Goldberg (2021) that empirical evidence such as the one discussed in Benkler et al. (2018) suggests that we should *not* group irresponsible journalism such as Breitbart and Fox News together with other more responsible media. At a minimum, to be rightly grouped within the relevant category, they should “follow the norms or processes of professional journalistic objectivity” (14). Our recipe simply requires that people consume more news from sources of this kind. But while just Fox News *claims* to be doing this, there is ample empirical evidence that *neither* news source, in fact, does. This is why the undeniable fact that these irresponsible media houses contribute to circulating misinformation does not jeopardize the prospects of our proposed therapy.

The second objection pressures us into refining the target of our suggestion. Suppose someone believes that all mainstream media routinely hide the truth to the public in the pursuit of some political

⁶ Recent political events reveal that this problem has reached new heights, as many U.S. news anchors have called out and interrupted former U.S. President, D. Trump, who was addressing Americans from the White House and claiming victory in the presidential election while millions of votes still had to be counted.

agenda. When “independent” sources report P and no mainstream media does, this person would arguably *fail* to have a reason to reject P as fake news. Since they believe that the traditional media *omit* unsettling truths on purpose, this person would more likely be rationally required to believe P and perhaps—somewhat perversely, from an epistemological point of view—increase her confidence in the media conspiracy. Crucially, however, this and similar cases are built around subjects whose informational situation is *so* pathological that their doxastic conduct deserves a different diagnosis and requires separate treatment, as we shall suggest in §4. The present analysis is only concerned with the doxastic conduct of *ordinary* people, that is, subjects who may suffer from limited or bent access to information but who have an otherwise functional conception of the social reality. These subjects, for instance, realize that traditional news outlets wouldn’t easily fail to report highly sensationalist information if they believed it was true news. So, when they realize that mainstream media systematically fail to cover a news report, they normally take it as an indication that the media judge it to be fake. Once the real target of our analysis is elucidated, the case of the paranoid conspiracist, albeit realistic, no longer menaces its plausibility.

As for the third objection, assessing the extent to which ordinary social media users consume news from traditional media is an open empirical question.⁷ No matter what verdict future data will provide, it deserves emphasis that our proposal should be intended as the comparative claim that the *more* news social media users consume from traditional media, the *less* fake news they can be expected to end up believing under the influence of their cognitive biases. This claim, in virtue of its comparative nature, would arguably remain undefeated even if it turned out that our recipe *is* currently being implemented to some degree. Our response, in this case, would be that to further reduce fake news consumption by the relevant subjects, the same recipe should be implemented *even more*.

⁷ According to a recent report by the Pew Research Center (“Americans Who Mainly Get Their News on Social Media Are Less Engaged, Less Knowledgeable”, July 2020, available at <https://www.journalism.org/2020/07/30/americans-who-mainly-get-their-news-on-social-media-are-less-engaged-less-knowledgeable/>), the Americans that primarily rely on social media for news are less likely to get things right about politics and other social phenomena. Although data such as these do not settle the empirical question looming large, they certainly *do not disprove* our main suggestion, according to which a more varied information diet would tend to reduce the likelihood of being fed with misinformation.

4. CONCLUSIVE REMARKS

This paper has attempted to defend educational strategies to counteract fake news consumption against the claim that they are poorly motivated because social media users cannot be criticized for their doxastic conduct. In particular, we have addressed two variants of this claim. In Rini's version, social media users cannot be criticized for believing in fake news because, in many cases, this is the result of an exercise of virtuous epistemic partisanship. In Millar's version, social media users are fully excused for consuming fake news, in that they mostly do it under the influence of psychological and environmental factors that they could not (be requested to) control. Against Rini, we have observed that most fake news consumed online is not within the range of application of the intellectual virtue of epistemic partisanship. Against Millar, we have argued that his analysis underestimates the advantages of having a more varied information diet and neglects the beneficial effect of *not finding* reports of fake news on mainstream media, which would have most likely reported it if ad been true.

Our critical analysis of these views has provided us with a constructive moral: namely, that helping social media users to have a more varied information diet amounts to a fundamental and practically feasible educational remedy. If we are right, educating social media users about the epistemic benefits of a varied information diet appears to be a more promising strategy than merely carving social media algorithms to ensure that they get exposed to a variety of information sources. For, while structural approaches can change the architecture of online information environments unbeknownst to their users, educational remedies provide them with the resources to understand the problems affecting these environments and manage their online epistemic conduct more responsibly.

In conclusion, let us go back to the remark that the educational therapy we have proposed is meant to be beneficial to *ordinary* social media users (see §3) and add that, by contrast, it is not going to bring about the desired effects when it comes to social media users trapped in what Nguyen (2020) calls *echo chambers*. Echo chambers are social epistemic structures that distribute epistemic credit

asymmetrically between insiders and outsiders by overinflating the epistemic trustworthiness of the former and overdeflating that of the latter irrespective of their actual epistemic worth and reliability.⁸ Epistemic agents within echo chambers are taught to dismiss any counterevidence they are exposed to as a corruption attempt on the part of outsiders. For this reason, asking them to widen their information diet has no prospects to reduce their proneness to consuming fake news. Quite to the contrary, it is likely to backfire, that is, to discredit contrary opinions and increase inner trust among the members of an echo chamber.

Our proposed educational remedy does not address the epistemic needs of social media users trapped within echo chambers, as they seem to pose problems of their own and therefore require a separate treatment. It will be enough, for the moment, if we have managed to indicate a viable way to ameliorate the online epistemic conduct of ordinary agents, who have to counteract the combined effects of cognitive biases and social media algorithms. To resort to a different metaphor suggested by Nguyen (2020), these agents are inserted within bubbles, where the information is impeded from circulating in a free manner. If what we have said in this paper is on the right track, these are bubbles that social media users can be easily taught how to pop.⁹

REFERENCES

Alston, W. (1988). The Deontological Conception of Epistemic Justification. *Philosophical Perspectives* 2: 257-299.

⁸ Nguyen distinguishes echo chambers from *epistemic bubbles*, where the latter are structures that impede the free circulation of information by omitting relevant sources. As we have seen, social media such as Facebook or Twitter are environments mostly populated by like-minded people and are sustained by algorithms that select the information that is allowed to circulate. Both factors make these virtual environments good examples of epistemic bubbles in Nguyen's sense.

⁹ Earlier versions of this paper have been presented at the Invited Speakers Seminar Series at the Department of Christian Philosophy, University of Innsbruck; the LanCog Seminar, University of Lisbon; and the workshop "Cognitive Distortions and Democratic Failures: The Impact of Fake News on Democratic Politics" at the University of Piemonte Orientale. We are grateful to the audience for helpful discussion. We also thank two anonymous reviewers of this journal for their comments on an earlier version of this paper. Michel Croce's work on this paper has been carried out at the University College Dublin and has received funding from the European Union's Horizon 2020 research and innovation programme for the project Policy, Expertise, and Trust in Action (PErTiA) under grant agreement No 870883. Michel Croce is currently based at Logos, University of Barcelona, where he works as a Beatriu de Pinós Fellow.

- Benkler, Y., Faris, R., Roberts, H. (2018). *Network Propaganda: Manipulation, Disinformation, and Radicalization in American Politics*. Oxford: Oxford University Press.
- Bozdag, E. (2013). Bias in Algorithmic Filtering and Personalization. *Ethics & Information Technology* 15:209-227.
- Croce, M., Piazza, T. (2021). Misinformation and Intentional Deception: A Novel Account of Fake News. In N. Snow, M. S. Vaccarezza (eds.), *Virtues, Democracy, and Online Media: Ethical and Epistemic Issues*. London: Routledge, pp. 49-63.
- Dentith, M. R. X. (2017). The Problem of Fake News. *Public Reason* 8(1-2): 65-79.
- Fazio, L. K., Brashier, N. M., Payne, B. K., Marsh, E. J. (2015). Knowledge Does Not Protect Against Illusory Truth. *Journal of Experimental Psychology General* 144(5): 993-1002.
- Fricker, M. (2007). *Epistemic Injustice: Power & the Ethics of Knowing*. Oxford: Oxford University Press.
- Gelfert, A. (2018). Fake News: A Definition. *Informal Logic* 38(1): 84–117.
- Goldberg, S. (2021). Fake News and Epistemic Rot: Or, Why We Are All in This Together. In T. Grundmann, A. K. Flowerree, & S. Bernecker (eds.), *The Epistemology of Fake News*. Oxford: Oxford University Press, pp. 265-285.
- Habgood-Coote, J. (2018). Stop Talking About Fake News. *Inquiry* 62(9-10): 1033-1065.
- Heersmink, R. (2018). A Virtue Epistemology of the Internet: Search Engines, Intellectual Virtues, and Education. *Social Epistemology* 32(1): 1-12.
- Jaster, R., Lanius, D., (2018). What is Fake News?. *Versus* 2(127): 207–227.
- Lazer, D. M. J., Baum, M. A., Benkler, Y., Berinski, A. J., Greenhill, K. M., Menczer, F., ... Schudson, M. (2018). The Science of Fake News. *Science* 359(6380): 1094-1096.
- Levy, N. (2017). The Bad News About Fake News. *Social Epistemology Review and Reply Collective* 6(8): 20-36.
- Levy, N. (2019). Due Deference to Denialism: Explaining Ordinary People's Rejection of Established Scientific Findings. *Synthese* 196: 313-327.

- Lewandowsky, S., Ecker, U. K. H., Seifert, C. M., Schwartz, N., Cook, J. (2012). Misinformation and Its Correction: Continued Influence and Successful Debiasing. *Psychological Science in the Public Interest* 13(3): 106-131.
- Millar, B. (2019). The Information Environment and Blameworthy Beliefs. *Social Epistemology* 33(6): 525-537.
- Mukerji, N. S. (2018). What is Fake News?. *Ergo: An Open Access Journal of Philosophy* 5(35): 923-946.
- Nichols, T. (2017). *The Death of Expertise: The Campaign Against Established Knowledge and Why It Matters*. New York: Oxford University Press.
- Nickerson, R. S. (1998). Confirmation Bias: A Ubiquitous Phenomenon in Many Guises. *Review of General Psychology* 2(2): 175-220.
- Nguyen, C. T. (2020). Echo Chambers and Epistemic Bubbles. *Episteme* 17(2): 141-161.
- Pariser, E. (2011). *The Filter Bubble: What The Internet Is Hiding from You*. London: Penguin Press.
- Pepp, J., Michaelson, E., Sterken, R. K. (2019). What's News About Fake News? *Journal of Ethics and Social Philosophy* 16(2): 67-94.
- Pritchard, D. (2021). Good News, Bad News, Fake News. In S. Bernecker, A. Flowerree, T. Grundmann (eds.), *The Epistemology of Fake News*, Oxford: Oxford University Press, pp. 46-67.
- Rini, R. (2017). Fake News and Partisan Epistemology. *Kennedy Institute of Ethics Journal* 27(S2):E43—E64.
- Smart, P. R. (2018). (Fake?) News Alert: Intellectual Virtues Required for Online Knowledge!. *Social Epistemology Review & Reply Collective* 7(2):45-55.
- Van Bavel, J. J., Pereira, A. (2018). The Partisan Brain: An Identity-based Model of Political Belief. *Trends in Cognitive Sciences* 22(3): 213-224.
- Vosoughi, S., Roy, D., and Aral, S. (2018). The Spread of True and False News Online. *Science* 359(6380): 1146-1151.

Worsnip, A. (2019). The Obligation to Diversify One's Sources: Against Epistemic Partisanship in the Consumption of News Media. In C. Fox & J. Saunders (eds.), *Media Ethics: Free Speech and the Requirements of Democracy*, London: Routledge, pp. 240-264.