Fake News and Epistemic Vice: Combating a Uniquely Noxious Market

Megan Fritts and Frank Cabrera

1.   Introduction

The topic of fake news has received increased attention from philosophers since the term became a favorite of politicians (Habgood-Coote 2016; Dentith 2016). Notably missing from the conversation, however, is a discussion of fake news and conspiracy theory media as a *market*. This paper will take as its starting point the account of noxious markets put forward by Debra Satz (2010) and will argue that there is a *pro tanto* moral reason to restrict the market for fake news. Specifically, we begin with Satz's argument that restricting a market may be required when i) that market inhibits citizens from being able to stand in an equal relationship with one another, and ii) this problem cannot be solved without such direct restrictions. Our own argument then proceeds in three parts: first, we argue that the market for fake news fits Satz's description of a *noxious* market; second, we argue against explanations of the proliferation of fake news that are couched in terms of "epistemic vice", and likewise argue against prescribing critical thinking education as a solution to the problem; finally, we conclude that, in the absence of other solutions to mitigate the noxious effects of the fake news market, we have a *pro tanto* moral reason to impose restrictions on this market. At the end of the paper, we consider one proposal to regulate the fake news market, which involves making social media outlets potentially liable in civil court for damages caused by the fake news hosted on their websites.

2.   Noxious Markets

In chapter 4 of *Why Some Things Should Not Be for Sale*: *The Moral Limits of Markets*, Debra Satz discusses what she calls "noxious markets." Satz is quick to disabuse her readers of the notion that hers will constitute a critique of markets in general, or of market-based economic systems; rather, she is concerned with "differing characteristics of very particular market exchanges: in human body parts, child labor, toxic waste, sex, and life-saving medicines" (2010: 91). She is, perhaps, most well-known for her discussions of kidney markets and markets for sex-work. Broadly, Satz seeks to identify categories of market-exchange that prey on the weak, exploit the ignorant, and cause widespread societal harm for reasons *inherent* to the type of exchange itself.

To help us identify such markets, Satz (2010: 94-8) lists four dimensions along which a market may be noxious:

(1) **Individual Harms:** the market produces extremely harmful outcomes for individual participants.

(2) **Societal Harms:** the market is extremely harmful to the society.

(3) **Weak Agency:** the market involves or requires very weak or asymmetric knowledge or agency on the part of participants.

(4) **Vulnerability:** the market reveals the deep vulnerabilities of one of the transacting participants.

Satz rightly recognizes that nearly all markets have one or more of these traits to some degree; that is, nearly all markets will be slightly noxious along at least one of these dimensions. Some markets which may be noxious along some of these dimensions we, nevertheless, may not wish to eliminate or restrict—for instance, tobacco products or pornography. We might not favor elimination or restriction of these markets because we believe the harm caused is only so caused because of a sort of clear-eyed willingness on the part of participating agents. Moreover, purely moralistic and paternalistic market restrictions go against the spirit of political liberalism, are difficult or impossible to implement, and have historically been used for oppression or to gain political advantage. Even when the proposed restriction is not moralistic, we might oppose legal measures to eliminate a market because we think, although the effects of the market are currently quite bad, these bad effects are contingent on other environmental factors and are therefore preventable (Nussbaum 1998).

Some markets, however, seem intrinsically harmful, and no amount of societal engineering or information awareness can prevent their ill effects. To identify such inherently noxious markets, Satz argues that we must find out which markets are likely, due to their *nature*, to be very harmful along one or more of these dimensions. As we will argue, the market for fake news displays noxiousness along *all four dimensions* in ways that are inherent to market, and therefore cannot be prevented without significant market restrictions.

3.   The Problem of Fake News

In this paper, we will be assuming the definition of fake news put forth by Allcott and Gentzkow (2017), according to which "fake news" picks out misleading articles that are "intentionally and verifiably false", including "articles that originate on satirical websites but could be misunderstood

as factual" (213). It is not, however, essential to our argument that our reader adopt this exact definition of fake news. Many other popular definitions are similar enough to do the trick. Gelfert (2018), for example, defines "fake news" as "the deliberate presentation of (typically) false or misleading claims as news, where the claims are misleading by design" (85-86); Rini's (2017) definition is similar to Gelfert's, but adds the additional clause that fake news stories mimic the appearance of real news stories, and are created with the goal of widespread transmission and deception. Allcott and Gentzkow's differs meaningfully from these definitions only in allowing satire to occasionally count as fake news. Given the recent incidents of politically charged satire being reposted by politicians as real news, we find this appropriate.[1]

In August of 2017, Facebook was in the news for the hundreds of fake news articles that had suddenly begun to appear on Facebook users' newsfeeds over the past year. The articles on Facebook users' newsfeeds were often entirely false or involved the conspiratorial demonization of one political candidate or another. Not only were these stories being read, but they were also being reposted, presumably because users thought they were real. Nearly half of the website traffic these fake news sites enjoyed came directly from social media sites—most often Facebook. The cause of immediate concern was the circulating fear that fake news articles, advertised and accessed through Facebook, had illicitly influenced the results of the 2016 US presidential election.[2] Immediately preceding the 2016 election, the top fake news stories, such as the headline from wtoe5news.com, which claimed that Trump had been endorsed by Pope Francis, were more widely read and shared than the most popular mainstream news stories (Silverman 2016). The societal implications of these numbers are obvious. With 62% of adults receiving their news from social media sites (Gottfried and Shearer 2016), and with social media sharing of top fake news stories far exceeding the sharing of top mainstream stories, the concern about undue political influence seems more than reasonable.

How these fake news stories have managed to garner so many clicks and views is not entirely clear, but a plausible answer is that the complex and ever-evolving algorithms of social media newsfeeds (especially Facebook) have the potential to give users the illusion that their virtual peer group is ideologically homogenous (Saez-Trumper et al. 2013; An et al. 2014).[3] This seeming

---

[1] https://www.usatoday.com/story/news/politics/2020/10/16/donald-trump-retweets-babylon-bee-attack-joe-biden-twitter/3675923001/

[2] In this paper, we largely focus on the fake news markets—and the harms of those markets—as they manifest in the United States. However, we believe that many of our arguments can generalize to other markets for fake news.

[3] For a detailed exposition of how these algorithms work, see Burr et al. (2018). As Burr et al. (2018) argue, these systems have become so sophisticated that it is useful to model them as autonomous agents attempting

homogeneity means that you are more likely to see (or, at least, pay attention to) posts with which you are already inclined to agree. For example, in a detailed study of the "anti-vaccine" and "pro-vaccine" communities on Twitter, Sullivan et al. (2020a) demonstrated that both communities are extremely isolated from each other, manifesting "a strong bias toward retweeting information from accounts that have the same viewpoint" (45). To borrow a distinction from Nguyen (2018), these "epistemic bubbles" are not on their own hostile to outside views and contrary evidence. However, the presence of fake news within these bubbles, which villainizes opposing views and fosters a sense of distrust of all sources of outside evidence, creates "echo chambers"; once shared, these inflammatory posts tend to spread like wildfire, and it becomes harder to convince readers that these stories are fake. Even if these false headlines are debunked by some other more reliable outlet, there is little assurance that those who read the original fake news headline will be exposed to the truth about the matter. As Mason et al. (2018: 5) point out: "Corrections or refutations are also now less likely to penetrate deeply into public consciousness partly because of ideological segmentation (meaning that those who would most likely be enlightened by the information are least likely to receive it)." As a result, such extreme isolation and homogeneity can leave the social networks of agents "epistemically vulnerable", as agents will be deprived of those ingredients that make for a healthy testimonial network: multiple, independent, and *diverse* sources (Sullivan 2020b: 6). In the case of vaccines, epistemic vulnerability can have dire consequences, as agents can be deprived of the expert advice of medical professionals.

The dynamics by which fake news spreads are similar to the processes of online self-radicalization, which have been investigated and contextualized by Alfano et al. (2018) as a kind of "technological seduction." Devotees of certain media outlets can have their worldviews influenced and can potentially be driven toward extremist views. This can happen in a "top-down" fashion, via deliberate editorial decisions, e.g. giving undue focus to certain news stories, employing sensationalistic rhetoric, categorizing articles with biased tags, etc., or in a "bottom-up" fashion, via the content recommendations generated for users by predictive analytics algorithms.[4] As Alfano et al. remark, "the consumer of (for instance) Breitbart will have her experiences structured in a way that may naturally

---

to maximize their own utility and possessed of many strategies to influence the behavior of human agents, e.g., coercion, nudging, trading, etc.

[4] See Alfano et al. (2020) for a detailed study of the YouTube recommender algorithm, in which the authors demonstrate that "there exists a robust pathway from certain seemingly anodyne topics to conspiracy theories via the recommender system" (3). Plausibly, the YouTube recommender systems contribute to the receptivity of fake news.

give rise to distinctive biases" (2018: 305). Plausibly, such biases will make its readers much more susceptible to certain fake news stories, *even if* Breitbart reports only veridical headlines.

Fake News as a Market

As we have suggested, the proliferation of fake news on social media has a complex causal profile, making it a complicated and multi-faceted market; yet, despite this complexity, it is clear that *economic transactions* are increasing the demand for fake news headlines. The structure of the fake news market can be understood as consisting of two different kinds of transactions—the first between advertisers and fake news creators, and the second between social media users and social media outlets.

Regarding the first type of transaction, fake news creators make the bulk of their money from advertising. Some owners of prominent fake news sites, such as Jestin Coler of the *Denver Guardian*, have claimed to make over $10,000 per month through selling advertisement spots to wealthy corporations like Google.[5] Fake news sites then utilize the algorithmic features of social media platforms that allow content shared by some users to appear on the newsfeeds of their like-minded peers. Fake news articles are dropped into social media "group pages" where individual users click, read, and share the link on their own page and to their peers. Social media sites are incentivized to permit fake news articles on their platform, because doing so allows *them* to collect valuable information on their users based on who interacts with these articles (more on this below). Here, money is made through advertising, with social media sites facilitating website traffic.

The second type of transaction involves the actual consumers of fake news. In this case, no *money* changes hands—rather, individuals allow the collection of their personal information, currently the most value commodity on Earth[6], in exchange for the ability to use social media platforms. Social media sites have no problem providing their services for "free", due to the amount of money they make on the data they—as well as fake news websites themselves—are able to collect from their users. Such data, as mentioned above, is then use in targeted advertising. Data on users' political affiliations is especially valuable, as demonstrated in the notorious Cambridge Analytica data breach[7]. Everything Facebook users do on its platform, including the consumption of real and fake news, is "paid for"

---

[5] See: https://www.cbsnews.com/news/whats-fake-news-60-minutes-producers-investigate/

[6] See: https://www.economist.com/leaders/2017/05/06/the-worlds-most-valuable-resource-is-no-longer-oil-but-data

[7] For an account of the Cambridge Analytica breach, see: https://www.nytimes.com/2018/04/04/us/politics/cambridge-analytica-scandal-fallout.html.

with their data. And legally-speaking, most of these transactions are voluntary and binding, usually described in prohibitively dense virtual "Terms and Conditions" agreements.

The Fake News Market as a Noxious Market

An examination of the market for fake news shows that it exhibits noxiousness along the first and second dimensions identified by Satz: the market causes great harm both to individuals and to societies in which fake news is prevalent. In addition to the pernicious effects of fake news on the 2016 election, which we have already discussed above, another societal harm of fake news is that it has the potential to lead to violence. Possibly the first time in recent history that fake news was seen as a genuine danger was December 4th, 2016—the day that Edgar Maddison Welch entered a restaurant with an AR-15, intent on finding and killing a ring of child sex traffickers. Of course, he found none, and no physical harm ultimately came to him or anyone inside the restaurant. But the origin of his belief—what has come to be known as "Pizzagate"—was a slew of fake news stories. These stories originated on an anonymous Twitter thread before being amplified by "infotainment" media pundits like Alex Jones. The claims were as follows: the police had discovered, in the classified emails found on Hillary Clinton's computer, the existence of a child sex trafficking ring being run out of a pizza shop called Comet Ping Pong. The details and rumors multiplied from there. Welch's consumption of this fake news had convinced him of the truth of these stories and, infuriated at the government's unwillingness to take immediate action, he decided to take matters into his own hands.[8]

More recently, fake news has likely played a significant role in undermining public health measures in the United States during the 2020 COVID-19 pandemic. The most striking example is the reception of the short conspiratorial film called "Plandemic", a pseudo-documentary described by the *New York Times* as "a slickly produced narration that wrongly claimed a shadowy cabal of elites was using the virus and a potential vaccine to profit and gain power."[9] The popularity of Plandemic eclipsed other pieces of popular news and entertainment, garnering over 2.3 million social media interactions in the first week and a half of its release. Pew Research polls conducted after the release of the film showed that at least 25% of Americans think the hypothesis that the COVID-19 pandemic is a planned conspiracy to be "probably true."[10]

---

[8] For an overview of other modern and historical instances of fake news, see O'Connor and Weatherall (2019).
[9] www.nytimes.com/2020/05/20/technology/plandemic-movie-youtube-facebook-coronavirus
[10] www.pewresearch.org/fact-tank/2020/07/24/a-look-at-the-americans-who-believe-there-is-some-truth-to-the-conspiracy-theory-that-covid-19-was-planned/

To be sure, it is underdetermined how many of these people were convinced by the movie. It is entirely possible that the film preyed on those with strong pre-existing beliefs about medical conspiracies—beliefs perhaps influenced by other forms of media, including the Twitter feed of former president Trump himself. As Alfano et al. (2020: 19) suggest, it might be "that users who are already inclined to accept conspiracy theories fill in the details with content from YouTube, not that YouTube takes non-conspiracy theorizers and turns them into conspiracy theorizers." Relatedly, Benkler et al. (2018: 233) discuss such "priming media", of the likes of Alex Jones, Sean Hannity, or the Drudge Report, designed to produce "outrage" in their viewers. While such media may not fall under our definition of "fake news" in every instance, it is crucial to note the importance of the "cathartic" effect of outrage media on the eventual acceptance of entirely fabricated headlines.

In addition, other aspects of conspiracy theory media have created crucial precedents for the success of Plandemic. Specifically, fake news involving vaccine safety has been on a steep incline over the last decade and shows no signs of slowing down.[11] A recent article in *Nature* named "viral misinformation" as the biggest risk to health during a pandemic (Larson 2018).[12] Additionally, fake news articles regarding the purported COVID-19 "cure" espoused by the Plandemic film—Hydroxychloroquine[13]—began circulating early in March 2020, before the release of the film. And while the full effects of this fake news on public health measures are still unknown, one truly noxious effect of media like Plandemic is the way the particular claims of the fake news media eliminate the possibility of outside disconfirmation or critique.

To help us understand how this type of fake news can sometimes be impervious to correction from disconfirming sources, Benkler et al. describe a "feedback loop" that exists between three social spheres: politicians, competing media sources, and the public, all of whom have self-serving incentives when it comes to political media narratives (2018: 77-78). Media sources, in their competition for viewers, find that their audience loyalty increases when they produce "identity-confirming news." Politicians, and those in service to a particular political agenda, deliver identity-confirming narratives to the public for the same reason that the media does—it increases the loyalty and enthusiasm of their voting constituents. These narratives provide more fodder for the competing media sources, thereby perpetuating the cycle.

---

[11] https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6122668/; see also: https://www.psychologytoday.com/us/blog/psych-unseen/201902/antivaxxers-and-the-plague-science-denial
[12] See also Mercier (2020: 201).
[13] A possible treatment for Sars-CoV-1, which was shown to be ineffective for treatment of COVID-19.

As we discussed in section 2, the fact that a market harms individual people, or even whole societies, is not on its own enough to generate an all-things-considered reason to close off that market. Satz agrees with Nussbaum that markets ought not to be closed off if their noxious effects can be mitigated in other ways. But Satz also argues that if markets undermine citizens' ability to interact equally—and if the problems with these markets *cannot* be overcome through other measures—then it may be necessary to prevent citizens from becoming transacting parties in the market exchange (2010: 104). In our context, that would mean curtailing the ability of citizens to produce or read fake news, thereby preventing them from being participants in the market. Again, this may only be justified if there are no other reasonable measures that could be taken to mitigate the noxious effects of the market. In the following section, we will argue that the most common suggestions for reducing the harmful effects of the fake news market can be reasonably expected to fail.

4.   Epistemic Vice and the Critical Thinking Response

A natural response to the recent proliferation of fake news and conspiracy theory media is that we ought to further emphasize "critical thinking" (e.g., McIntyre 2018: 59-60).[14] One might hypothesize that the epistemic failures manifest in socio-political discourse are the result of insufficient attention to established forms of reasoning and the methods of scientific inference. These are the sorts of critical thinking skills that are covered in standard introductory courses offered by almost every university philosophy department. Such courses are typically titled "Critical Thinking", "Reason and Argument", "Logic and Reasoning", etc. This response gains support from recent discussions in "vice epistemology" (e.g., Cassam 2016), which often connect belief in debunked conspiracy theories to the presence of epistemic (or "intellectual") vices, such as wishful thinking, gullibility, etc. Call any proponent of educational interventions to combat the presence of epistemic vices, broadly construed to include dispositions to employ fallacious argument forms but also *character traits* like gullibility, the *Educationist*. [15] According to the Educationist, just as our primary and middle schools stress the

---

[14] See also: "In a Fake News Climate, Critical Thinking Skills Are More Important Than Ever", in the Chronical of Higher Education. https://www.chronicle.com/paid-article/in-a-fake-news-climate-critic/36

[15] Since we are worried about people believing wildly false headlines, in our discussion of the Educationist we will focus primarily on the "reliabilist" account of epistemic virtues/vices (e.g., Sosa 1991). This account analyzes epistemic virtues as dispositions to rely on truth-conducive argument forms. Popular proposals to increase critical thinking education should be viewed as a call to promote reliabilist virtues. We will, however, touch upon the "responsibilist" account of epistemic virtues/vices (e.g., Zagzebski 1996) at the end of this section. This account analyzes epistemic virtues as praiseworthy intellectual *character traits* marked by an intrinsic motivation for truth.

importance of liberal democratic values, so too the value of virtuous inquiry ought to be stressed in our secondary and tertiary schools. Since the state is already committed to emphasizing certain political values, and since a functioning democracy depends on epistemically virtuous inquirers, this intervention would be wholly justified.

Of course, the success of the Educationist's proposal depends crucially on whether an epistemic vice explanation of the current situation regarding fake news *really* is the best explanation; but, it is far from clear that this is the case. First, it is doubtful that the increased acceptance of fake news is best explained by an appeal to epistemically vicious traits, habits, or dispositions. Insofar as one thinks that the current epistemic situation is uniquely worrisome, as we do, then the vice explanation requires commitment to the position that overall epistemic virtue has dropped precipitously in recent years. It would be extraordinarily surprising, however, if the epistemic virtue of the discussants in public debates had suddenly declined in such a way.

Second, and most important, the problem with the vice explanation of the current epistemic situation is that it is far from clear that those who believe in outlandish conspiracy theories and fake news are making any *formal* or *logical* mistake. That is, it is far from clear that those who believe in fake news are employing any invalid or fallacious form of inference. If this is right, then proposed interventions to better educate the public about the proper rules of inductive and deductive reasoning will prove causally inefficacious. In any given case, it is of course difficult to know exactly on what basis a person comes to believe whatever demonstrably false claim is at issue. But the important point for our purposes is that the believer of some conspiracy theory or fake news headline *can be* modelled as respecting the formal rules of inductive and deductive inference. That is, these methods do not in themselves *rule out* belief in conspiracy theories and fake news headlines.

To see this, let's consider the most prominent model of scientific reasoning among philosophers of science, namely Bayesian confirmation theory (Douven 2011). According to Bayesianism, rational agents can be thought of as having "degrees of belief" or "credences" with respect to propositions, a notion which corresponds intuitively to how strongly the agent believes the proposition. For an agent to be rational, her "credence function", or "personal probability function" (Pr-function) must be constrained so as to obey the axioms of the probability calculus, and moreover, the agent must update over time in light of new evidence by the following rule:

> *Conditionalization*: For any hypothesis H and any body of evidence E, upon learning that E is true, the agent's doxastic state ought to be such that $Pr_{New}(H)=Pr_{Old}(H|E)$.

According to Conditionalization, the agent's old "posterior probability" of H given E, i.e. $Pr_{Old}(H|E)$ becomes the agent's new "prior probability" of H, i.e. $Pr_{New}(H)$, once the agent learns that E is the case. Bayesianism thus gets off the ground only if there is some initial prior probability distribution that was *not* itself the result of a previous instance of Conditionalization. It is because of this feature that we can often model those who are disposed to believe in fake news as rational Bayesian inquirers.

Let's consider a simple example. Before the 2016 election, one of the most shared fake news stories was "Pope Francis shocks world, endorses Donald Trump for president", which came from the now defunct website WTOE5 News. Let W be "WTOE5 reports that X", where X is some proposition, and let E be "Pope Francis endorses Donald Trump for president." For an agent to rationally have a high credence that Pope Francis endorses Donald Trump for president, it is sufficient that the part of the agent's credence function that bears on this issue looks like the following:

E	W[E]	Pr(-)	
T	T	.45	=Pr(E&W[E])
T	F	.05	=Pr(E&~W[E])
F	T	.05	=Pr(~E&W[E])
F	F	.45	=Pr(~E&~W[E])

Since the values of the conjunctive probabilities add up to 1, this probability distribution is coherent. From the law of total probability, the agent's unconditional, prior probabilities can be computed[16]:

Pr(E)=.5

Pr(~E)=.5

Pr(W[E])=.5

Pr(~W[E])=.5

As it happens, the agent is exactly indifferent between both of our relevant propositions and their negations. But, given our truth table, we can compute the probability of E conditional on W[E]:

$$Pr(E|W[E]) = \frac{Pr(E\&W[E])}{P(W[E])} = \frac{.45}{.5} = .9$$

---

[16] So, for example, Pr(E)=Pr(E&W[E])+Pr(E&~W[E]) = .5, etc.

Since Pr(E)=.5, and since Pr(E|W[E])=.9, it follows that Pr(E|W[E])>Pr(E), and thus according to the standard Bayesian account of confirmation, that W[E] confirms E. Most significant, given this probability function, if the agent learns that W[E] is the case, and then comes to have a credence in E that is very high—in this case .9—then she *need not* violate any Bayesian norm. If the agent has the above credence function, then she is rationally obligated by Conditionalization to update her credence in E to .9, and would, plausibly, be rational to believe, full-stop, that E is true.

Now, of course, one has the right to ask the agent why she has the above initial probability distribution. For all the merits of the Bayesian framework, prior probabilities have long been a thorn in the Bayesian's side (Sober 2002). According to standard varieties of Bayesianism, provided the agent's initial Pr-function obeys the axioms of probability theory and the agent adjusts her credences in accordance with Conditionalization, then the agent's doxastic state is in good rational standing. This leaves open a wide variety of rationally permissible Pr-functions, including ones that seem intuitively rationally indefensible, such as those which would license inductive skepticism (Huemer 2009).

It is plausible that one who adopts something like the above prior probability distribution is antecedently skeptical of the "mainstream media" and puts more stock in "alternative news sources." In a discussion of the rationality of conspiracy theorizing, Levy (2019) argues that despite being "objectively irrational and subjectively irrational for most of us", "accepting conspiracy theories is, for many people, subjectively rational" (66). As Levy points out, many of those who accept bizarre conspiracy theories perceive themselves as marginalized by society, and thus "exhibit comparatively low levels of trust toward those they see as representing the powerful; accordingly, they will be less likely to accept the official story" (73). Instead, such individuals manifest "relatively high levels of trust toward *un*official sources: toward in-group members and toward others…working against or in opposition to the elites" (73). One might fault the conspiracy theorizer for the sources that they trust, but in assessing the reliability of some testimony, "to some degree, we have no choice but to assess both competence and benevolence [of the testifier] by our own lights", and so "it is unsurprising that we prefer testimony from people who resemble ourselves: from those who share our partisan leanings…and those who share our background beliefs" (70). [17]

In any case, what is clear is that why the agent adheres to this asymmetry about the reliability of news sources depends on investigating a complicated set of background beliefs that takes us far

---

[17] We are quite sympathetic to Levy's position, and our general discussion of fake news here complements his specific discussion about the rationality of conspiracy theorizing. Indeed, Levy even mentions, but does not explore in depth, the idea that "conspiracy theorists have unusual Bayesian priors" (2019: 66).

beyond the realm of inductive logic. From our perspective, such claims regarding the unreliability of the mainstream media are, in an objective sense, epistemically unjustified. But no amount of emphasis on proper inferential schemas will demonstrate this fact to those who believe otherwise. Plausibly, if we presented the believer of a fake news headline with the Bayesian formalism, she would retroactively see herself as endorsing the prior probability distribution that preserves the rationality of holding the belief at issue. Indeed, there is some empirical evidence to suggest that those with greater education are *more* prone to belief in fake news and conspiracy theories, as "being more knowledgeable provides more ammunition with which to counter unpalatable claims" (Levy 2017: 33).

It is worth mentioning that the problem that we have highlighted here is not one that merely affects Bayesianism. One might view the above example as sufficient reason for the Educationist to focus on a different account of scientific reasoning, such as Inference to the Best Explanation (IBE). Often, IBE is formalized as a four-step argument pattern (e.g. Psillos 2002) such as:

(i)   F is some fact or collection of facts

(ii)   Hypothesis H_1, if true, would explain F

(iii) H_1 is a better explanation of F than its competitors H_2, H_3,...,H_n

(iv)   Therefore, probably, H_1 is true

Here, H is upheld as rationally justified by showing that H would, if true, explain some set of facts better than its competitors. Whether H is the better explanation depends on how well H does with respect to certain "explanatory virtues", e.g., *simplicity*, *scope*, *fit with background knowledge*, etc.

An appeal to IBE will not be a panacea for the problem of fake news, however. Perhaps most troubling for the Educationist is that even IBE needs to rely on considerations of prior probabilities. As Lipton (2004) argues, any plausible version of IBE must be such that the pool of all possible explanations is narrowed by plausibility considerations (59-61). According to Lipton's "two-filter" model of IBE, these plausibility considerations are understood as a "non-explanatory notion of likeliness" (2004: 61), or in other words, as the prior probability of the explanation.[18] For this reason, even someone hostile to Bayesianism and friendly to IBE must contend with a version of the problem of the priors. But this means that we are back to where we started. The believer in fake news will view the veridicality of the fake news headline as the better explanation, once again because of substantive background beliefs, which ground the acceptance of a particular, deviant set of prior probabilities.

---

[18] Lipton (2004, Ch. 7) also argues that IBE should be embedded in the Bayesian framework, which is another reason that the problem of the priors is not just a problem for Bayesians.

The upshot of the foregoing analysis is this: rather than making any *formal* or *logical* error, the mistake that those who are disposed to believe in fake news headlines make is *material*, having more to do with the substantive, epistemic background conditions that lead to endorsing one initial probability distribution over another. In our view, the epistemic background conditions have been sufficiently polluted so as to render any would-be intervention to better educate the public about the proper rules of inductive and deductive reasoning causally inefficacious. As a result, the critical thinking response endorsed by the Educationist is likely to fail.[19]

Objection: Informal Fallacies and Media Literacy

One way in which the Educationist might respond to our critique above is to claim that our understanding of "critical thinking" is impoverished. Of course, says the objector, increased emphasis on critical thinking will be causally impotent with respect to the problem of fake news if all that you mean by "critical thinking" is training in deductive and inductive logic. However, the objector goes on, there is more to critical thinking than learning valid argument forms. In addition to formal argument structures, a proper critical thinking course ought to include instruction in basic "media literacy" and discussion of informal fallacies/cognitive biases. For instance, McIntyre (2018) connects the problem of fake news to wider concerns about a "post-truth" society and suggests that "one of the barriers to critical thinking is bathing in a constant stream of confirmation bias" (p 163). The objector might agree then with our diagnosis that the problem of fake news is *material* rather than *formal* but go on to insist that critical thinking courses ought to be constructed in such a way as to address this problem. And one way to do so is to include in the critical thinking curriculum modules that instruct students on how to identify good material for the construction of arguments, which would include instruction on how to distinguish reliable sources of information from unreliable ones.

This is a legitimate challenge, and we wish that this response on behalf of the Educationist were successful, but unfortunately, we are skeptical for a number of reasons. First, it needs mentioning that our conception of critical thinking is hardly non-standard. To be sure, it is also commonplace in critical thinking courses to discuss informal fallacies/cognitive biases, e.g., appeal to authority, the genetic fallacy, the post hoc fallacy, etc. However, recent research in argumentation theory has shown

---

[19] While our focus has been on individualistic epistemology, a similar problem arises if we shift to network epistemology. As Pallavicini et al. (2018) demonstrate, "groups of ideally rational Bayesian agents who inquire and communicate with one another about some proposition exhibit group polarization behavior under a broad range of circumstances."

that the way informal fallacies are typically treated in textbooks suffers from serious problems. As Boudry et al. (2015) have convincingly argued, the standard informal fallacies are often uncharitably regarded as deductive arguments, which are obviously invalid, but also which have very few instances in everyday life. But if the alleged fallacies are more charitably interpreted, it will become clear that they are in fact legitimate, defeasible forms of reasoning, whose success depends on highly contextual background factors or other domain-specific considerations.[20] In either case, educating students about informal fallacies won't be particularly useful, and in some ways the popular "theory of fallacies" is counter-productive to the goal of producing better thinkers. As Boudry et al. conclude: dismissing an argument as a fallacy "in everyday conversation is akin to *name calling*—occasionally effective for practical purposes, but rarely informative and basically unfair" (2015: 452; emphasis added).

Second, while we regard the suggestion that a critical thinking course should include a module focusing on media literacy as promising, again, we are skeptical about the feasibility and causal efficacy of such a proposal. Our major objection to this proposed module is that designing it in a way that will be perceived as non-partisan will prove quite an obstacle. To illustrate our point, consider following partial list of course goals from a course entitled "Media Literacy: Criticism & Analysis":

> Identify who created a particular media message, b) Recognize what the media maker wants us to believe or do, c) Name the 'tools of persuasion' used, d) Recognize bias, spin, misinformation and lies, e) Discover the part of the story that's not being told, f) Evaluate media messages based on our own experiences, beliefs and values.

In the abstract, these course goals sound like wonderful ideas; it would be a great boon to students if they could develop these skills. But now, consider also that this course is offered by a private secondary school called "Life Christian School" in Hillsboro, Oregon.[21] We are confident that the perception of this course by some of our readership has now radically changed in predictable ways.

In our highly politically polarized era, a media literacy course of this kind, with a focus on substantive *material*, rather than *formal* problems in reasoning, or a critical thinking course with a media literacy component, will be viewed with great suspicion. And to some degree, this suspicion is not unfounded. In an evaluation of media literacy efforts in Israel, it was concluded by Lemish and Lemish (1997), perhaps unsurprisingly, that policy-makers tended to perceive the media from the standpoint of, and propose media literacy interventions consistent with, their own political ideologies. While we are confident that a media literacy module *could* be taught in a non-partisan way, in this case it is

---

[20] See, for instance, a discussion of the "genetic fallacy" in which Sober makes this point (1994: 104-6).
[21] See https://canvas.instructure.com/courses/766248/assignments/syllabus for the full syllabus.

precisely the perception, and not reality, that matters. Those who would most benefit from media literacy will, in our estimation, be least likely to take or take seriously such a course, if the course is viewed as politically slanted. But it is just this outcome that is highly likely to occur.

Furthermore, and unfortunately for the Educationist, there exists little empirical data to support the hypothesis that media literacy interventions will be causally efficacious in combating fake news. According to a 2018 report titled "The Promises, Challenges and Futures of Media Literacy", drafted by the *Data & Society Research Institute* and published in the *Journal of Media Literacy Education*:

> In general, there is a lack of comprehensive evaluation data of media literacy efforts. Some research shows that media literacy efforts can have little-to-no impact for certain materials, or even produce harmful conditions of overconfidence (Bulger and Davison 2018: 1).

The report acknowledges that there is a "dearth of rigorous research pairing media literacy education with outcomes" (Bulger and Davison 2018: 10-11) and concludes, somewhat pessimistically, that "from an evidence perspective, there remains uncertainty around whether media literacy can be successful in preparing citizens to resist 'fake news' and disinformation" (11-12).

A recent observational study by Jones-Jang et al. (2019) on the relationship between media literacy and fake news casts further doubt on the above educational proposals. In order to test the hypothesis that media literacy improves one's ability to identify fake news, Jones-Jang et al. first distinguish four distinct, yet closely related types of literacies that have received discussion in the literature: a) media literacy, b) information literacy, c) news literacy, and d) digital literacy. First, *media literacy* is in this study defined as "the ability of a citizen to access, analyze, and produce information for specific outcomes" (Aufderheid 1993: 6). Among other things, media literacy defined narrowly in this way involves the ability to discern political or ideological biases of source material. Next, *information literacy* is defined by the Association for College and Research Libraries as an "intellectual framework for understanding, finding, evaluating, and using information" (ACRL, 2000). Among other things, information literacy involves the ability to identify and locate established knowledge, by means of, for instance, a library database. Third, *news literacy* is defined as "motivation to seek out news, the ability to find/identify/recognize news, the ability to critically evaluate and analyze news, and the ability to produce news" (Jones-Jang et al. 2019: 5). Among other things, news literacy typically involves a comprehension of the broader context in which news is produced, such as the fact that news organizations often have profit-motivations. Finally, there is *digital literacy*, which can be defined as the ability to "adapt constantly to new technologies" (Jones-Jang et al. 2019: 6). Among other things, digital literacy involves the ability to understand important online search-related terms such as *wiki* or

tagging. What Jones-Jang et al. sought to determine is if measures of these various forms of literacy predict the ability to distinguish between real news headlines and fake news headlines.

Sadly, the results of the study were far from encouraging. Of the four forms of literacy distinguished, *only* information literacy—which was measured by asking participants questions such as "Which of the data listed below are 'raw' unprocessed data?"—was positively associated with fake news recognition. Media literacy, as defined in this study, was *not* positively associated with detection of fake news. While one might be heartened by the fact that at least *one* form of literacy predicts the ability to discern fake news from real news, there are several important limitations to keep in mind. First, the effect size was rather small (regression coefficient =.119). Second, since the study was observational, it is difficult to know what sort of intervention would be causally efficacious in light of the findings of the study. As Jones-Jang et al. point out, "the findings do not reveal how individuals identify fake news stories" (2019: 14). Some may have identified the fake headlines by relying on prior knowledge, whereas others may have made educated guesses. Since information literacy involves abilities such as effectively using a library research database, it is unclear what causal connection, if any, there is between these skills and the ability to detect fake news headlines. Finally, in keeping with the theme of this section, Jones-Jang et al. caution, in spite of (or perhaps because of) their modest finding, that "an undue emphasis on audience education may oversimplify the issue… in programming effective literacy interventions, it is essential to recognize the limits of users' cognitive ability and resources" (2019: 14). In our assessment of the Educationist's response, we can't help but concur.

Can Teaching Virtuous Character Traits Help?

Thus far we have focused on proposals to increase critical thinking skills and media literacy, but one might think that the real educational intervention needed to combat the problem of fake news is one that improves people's *character*. Among virtue epistemologists, there are two ways to analyze epistemic virtues (Battaly 2016: 163). Whereas "virtue reliabilists" understand epistemic virtues as cognitive processes/dispositions that reliably lead to truth, e.g., deductive reasoning, perception, etc. "virtue responsibilists" understand epistemic virtues as praiseworthy intellectual character traits, e.g., open-mindedness, epistemic humility, etc. To possess these responsibilist virtues requires "caring about truth for its own sake" (Battaly 2016: 181). While we have focused primarily on the reliabilist virtues in our discussion of potential educational interventions, we could imagine the Educationist endorsing the teaching of responsibilist virtues in critical thinking courses *alongside* the teaching of reliabilist virtues (e.g., Battaly 2016; Heersmink 2017).

Although the intrinsic motivation for truth underlying the responsibilist virtues can plausibly contribute to the acquisition of true beliefs (Battaly 2016: 171), we have several concerns with this last suggestion for dealing with the problem of fake news. First, cultivating this motivation will prove difficult and time-consuming, as doing so requires "opportunities for practice across the curriculum and throughout our lives" (Battaly 2016: 181-2). Second, as Tanesini (2016: 524) points out there is a general problem with any attempt to teach the responsibilist virtues, as "those students who are the furthest away from intellectual virtue are precisely those who are less likely to pay attention." Similarly, Cassam (2015) has dubbed some particular epistemic vices "stealthy." Such vices, by their very nature, are hard to detect by someone who has them. A closed-minded person, for instance, is likely not *open* to the possibility that they have this epistemic vice. For these reasons, we are doubtful that focusing educational interventions on inculcating responsibilist virtues will prove any more efficacious than the suggestions geared more toward reliabilist virtues.

5.   Addressing Weak Epistemic Agency: Some Attempts and Failures

As we discussed in section 3, the salient features of much of the fake news market—the doubt cast on all other disconfirming sources—fulfils Satz's first and second criteria of noxious markets, by causing harm both to individuals and to society. Further, as we have shown in the preceding section, the market for fake news also exhibits noxiousness along the third and fourth of Satz's dimensions: it relies on and deepens weak agency and exploits the underlying epistemic vulnerabilities of the transacting parties. We consider the market for fake news to be a unique kind of noxious market, at least in comparison with most other noxious markets, because it relies on and deepens weak *epistemic* agency, and exploits (in particular) underlying *epistemic* vulnerabilities. That many of the harms are specifically epistemic in nature[22] is incredibly important for analyzing the noxiousness of the market— namely, it is the reason that standard educational solutions are unlikely to be effective.

The noxious effects of the fake news market run together in many ways: for instance, a community rendered collectively epistemically inculpable by a vastly-polluted evidence pool will likely be composed of people who have all been *individually harmed,* thereby manifesting the first trait of a noxious market. Individuals are likely to be harmed, not only epistemically, but also more literally, as the effects of the polluted evidence pool become more severe. For instance, individuals may find themselves in a society run by megalomaniacal politicians, in a world suffering severe environmental

---

[22] These epistemic harms may also constitute a violation of individuals' "epistemic rights" (Watson 2018).

degradation, in a nation dragged down by failed public health measures, or even in a family unable to communicate with each other about the most important issues. Likewise, a market that requires and deepens weak agency will almost certainly "reflect the underlying extreme vulnerabilities of one of the transacting parties" (Satz 2010: 87), because those who enter the fake news market already in a weak position (i.e., those who are already epistemically vulnerable, either by their evidence pool, peer group, or something else) will be much more susceptible to the agency-corroding effects of the market.

But if, as we have argued, the proposals of the Educationist are likely to prove ineffective, what else could be done about the market for fake news? Satz argues that a potential solution to weak agency is to increase public awareness: "[I]f weak agency is the problem with a particular market, then we may want to undertake measures that increase information" (2010: 104). Accordingly, as a response to the deluge of fake news stories on their social media platform, Facebook launched a highly publicized but short-lived campaign that attempted to "flag" stories that contained various markers of fake news. Users were given the ability to indicate that the article was fake news, by clicking a button located next to the headline of a posted news article. A tagged news story was then evaluated by third party "fact-checkers", and either taken down or left alone. Alfano et al. (2020: 19) suggest, in the case of YouTube, that perhaps problematic videos also ought to be flagged and reviewed by a human. A similar strategy is suggested by Regina Rini (2017), who argues that social media sites should allow users to flag, not only fake news sites, but individual users themselves.

Unfortunately, such measures will not be effective in markets like the one for fake news because, as we addressed in the previous section, the market for fake news relies on and exacerbates weak *epistemic* agency. More than being merely ineffective, Facebook's flagging plan actually *increased* partisan opposition to this campaign and resulted in more deeply entrenched beliefs in the claims of the disputed articles. Facebook product manager Tessa Lyons, in an official blog post announcing the cessation of Facebook's flagging campaign, admitted: "Academic research on correcting misinformation has shown that putting a strong image, like a red flag, next to an article may actually entrench deeply held beliefs – the opposite effect to what we intended."[23] One can easily imagine the same scenario playing out in instances of "user flagging", with those who believe fake news stories taking pride in their low Reputation Score and distrusting those with high Reputation Scores.

The inevitable result of Facebook's original campaign to flag fake news was the suspicion of purely partisan motivations. The fear was that the "triggers" for a news article being flagged as fake

---

[23] https://www.bbc.com/news/technology-42438750

boiled down purely to the presence of politically conservative terminology. Commenting on Facebook's campaign, Mark Epstein of the *National Review,* a conservative news website, claimed, "Conservatives have long accused these fact-checkers of liberal bias, and their new powers exacerbate these concerns…Determining what's a reputable outlet or fake news requires subjective decisions on controversial issues."[24] In an effort to ease the worries of political conservatives, Facebook hired the conservative-leaning news website *The Weekly Standard* as part of their fact-checking team. But this move backfired as well. "As was entirely predictable, *The Weekly Standard* is using its new authority as a Facebook fact-checker to label liberal arguments 'false'," groused *The New Republic.*[25]

6.   Some Conjecture about Potential Legal Solutions

If attempts to directly provide the public with accurate information are also helpless against the threat of minimized epistemic agency, then what else is left? Satz recognizes that some markets may be inherently noxious and may therefore require restriction or elimination. Satz notes that in these cases "our goal should be to curtail a particular noxious market, not to make it work better" (2010: 94). If there are no sufficiently countervailing reasons, then the *pro tanto* moral reason to limit the market for fake news because it is a noxious market becomes an *all-things-considered* reason. While we do not commit ourselves here to the conclusion that we have an all-things-considered reason to restrict fake news, we nonetheless find this possibility sufficiently plausible to be worth considering. On the assumption that the noxious nature of the market for fake news gives us a decisive reason to restrict it, then clearly some legal or regulatory response is in order. Fortunately, the feature of the fake news problem that makes direct extermination (practically) impossible is the very feature supplying the market with blood—social media platforms. The market for fake news, then, may be controlled or eliminated by directing legislation toward social media companies.

One way this may be done, which we find promising, is by repealing Section 230 of the Communications Decency Act (CDA). This section, which has been called the "most important law protecting Internet speech" (Electronic Frontier Foundation)[26], states that "No provider or user of an interactive computer service shall be treated as the publisher or speaker of any information provided by another information content provider." Under this section of the law, owners of social

---

[24]   See: https://www.nationalreview.com/2016/12/fake-news-facebook-new-york-times-hate-speech-socialmedia-political-diversity/

[25] See: https://newrepublic.com/minutes/151171/facebook-uses-weekly-standard-fact-check

[26] https://www.eff.org/issues/cda230

media sites are exempt from all legal responsibility for the content posted by users of their site—this includes libel and content encouraging violence or otherwise criminal activity. While expansive content moderation would in any case pose a "particularly gnarly challenge" (Benkler et al. 2018: 365), Section 230 of the CDA provides further disincentives for social media outlets to even attempt such a feat. As we discussed in section 3, because the spread of fake news increases the profits of both the producers of fake news and the sites on which the news is propagated, site owners currently have a large *financial incentive* to allow for all manner of misinformation. Repealing Section 230 of the CDA could do much to extinguish these perverse incentives, as the threat of imminent litigation would become a reality for website owners.

A concern with this proposal is Donald Trump's recent endorsement of it. Trump's endorsement followed the events of Twitter "censoring" many of his tweets alleging fraud in the 2020 presidential election. Such tweets often contained fake news or misinformation, which posed a serious threat to public well-being, and as a result warning labels were appended to these tweets. Trump appears to be in favor of the repeal of Section 230 because he believes it would allow users to sue site owners for censorship and defamation. Indeed, recently an executive order was issued seeking to make social media companies more vulnerable to censorship lawsuits.[27] Importantly though, repealing Section 230 of the CDA has bi-partisan support. In December of 2019, then-candidate Joe Biden endorsed repealing Section 230.[28] A year later, a top tech advisor in the Biden administration recommended "throwing out Section 230 and starting over."[29] The end of Section 230 may be on the horizon, but several potential drawbacks will need to be addressed in order for this measure to prove effective, all-things-considered, in curbing the market for fake news.

Regarding Trump's charges of censorship, there is little evidence that any lawsuits alleging free speech violations on the part of social media companies would carry any weight. Tech companies are private entities with First Amendment rights which have repeatedly held up in court.[30] An explosion of frivolous lawsuits alleging defamation is, however, a more worrisome problem. As Flick points out, "The 'specter of tort liability' remains the largest perceived threat by ICSPs [Interactive Computer Service Provider] of any new reshaping or reduction of Section 230 immunity by the CDA" (2017: 393). We suggest that any repeal of Section 230 must be accompanied by additional legal measures to

---

[27] https://www.whitehouse.gov/presidential-actions/executive-order-preventing-online-censorship/
[28] https://www.nytimes.com/interactive/2020/01/17/opinion/joe-biden-nytimes-interview.html
[29] https://www.cnbc.com/2020/12/02/biden-advisor-bruce-reed-hints-that-section-230-needs-reform.html
[30] https://www.theverge.com/2020/5/27/21272066/social-media-bias-laura-loomer-larry-klayman-twitter-google-facebook-loss

prevent over-litigiousness in civil court. And any future statutes, of course, must be carefully crafted to avoid violating the First Amendment rights of both individuals and corporations. There is a further concern that repealing Section 230 would be insufficient for curbing fake news, as the harms associated with fake news may be too distantly related to the social media platforms to count as legally actionable harms. As Flick suggests, such a statute would need to identify a certain class of "material, specific harm in narrowly construed contexts" and would require "specific knowledge of falsity by the author or publisher to be actionable" (2017: 396). A detailed examination of the proposal to repeal Section 230 is, of course, beyond our scope here; however, we believe, given the harms of the fake news market that we have identified, such a proposal is worth taking seriously. [31]

7.   Concluding Remarks

In this paper, we have aimed to discuss fake news as a market. This market is created and fed primarily by social media platforms which allow fake news articles to be target-advertised to the audiences that will be the most sympathetic to specific headlines. Furthermore, as we've argued, this market is a *noxious* market, leading to many harmful outcomes in society and damaging the epistemic agency of many of the transacting parties in the market. One popular response to the problem of fake news is to recommend increased education in critical thinking for the general public. Solving the problem of fake news, according to this response, involves ridding people of their epistemic vices— generally by way of expanding their critical thinking abilities. But, as we've argued, it is far from clear that the rise of belief in fake news headlines is best explained by appealing to epistemically vicious traits or dispositions. Rather than individual epistemic vices, the more plausible source of the problem is the pernicious background conditions in which agents are situated. Consequently, as we have shown, there is considerable reason to believe that increased emphasis on crucial thinking skills or media literacy will not serve to mitigate the noxious effects of the fake news market.

We conclude that, in the absence of other solutions to counteract the noxious effects of this market, we have a *pro tanto* moral reason to inhibit or eliminate this market. While an in-depth discussion of a legal-regulatory solution to the problem of fake news is outside the scope of this paper, we are particularly interested in proposals to amend or repeal Section 230 of the Communications Decency Act. The goal of such an amendment would be to hold site owners accountable for damages caused by the content found on their sites—either content posted by users or advertised to users. We

---

[31] For more extensive discussion of the legal challenges to curbing fake news, see Flick (2017).

briefly discussed some drawbacks to this proposal, which would require further legal measures to mitigate. These drawbacks include the potential for over-litigiousness toward website owners and the concern that the damages caused by fake news are too indirect for website owners to be held liable. However, we are hopeful that future work by philosophers, legal theorists, and policymakers could help to address these concerns. The possibility of social media outlets being held accountable in civil court may go a long way toward combating the noxious effects of the market for fake news.

Works Cited

Allcott, H. & Gentzkow, M. (2017). "Social Media and Fake News in the 2016 Election". *Journal of Economic Perspective*s. 3(2):211–236.

Alfano, M., Carter, J.A., & Cheong, M. (2018). "Technological Seduction and Self-Radicalization," *Journal of the American Philosophical Association* (3):298-322.

Alfano, M., Fard, A.E, Carter, J.A., Clutton, P. & Klein, C. (2020). "Technologically scaffolded atypical cognition: The case of YouTube's recommender system," *Synthese*:1-24.

An, J., Quercia, D., & Crowcroft, J. (2014). "Partisan Sharing: Facebook Evidence and Societal Consequences." In *Proceedings of the Second ACM Conference on Online Social Networks*. http://cosn.acm.org/2014/files/cosn033f-anAembTS.pdf

Association for College and Research Libraries. (2000). *Information literacy competency standards for higher education.* Chicago, IL: American Library Association. http://www.ala.org/acrl/standards/informationliteracycompetency.

Aufderheide, P. (1993). *Media literacy. A report of the National Leadership Conference on Media Literacy.* Washington, DC: Aspen Institute, Communications and Society Program.

Battaly, H. (2016). "Responsibilist Virtues in Reliabilist Classrooms." In J. Baehr (ed.), *Intellectual Virtues and Education: Essays in Applied Virtue Epistemology.* London: Routledge, 163–187.

Benkler, Y., Faris, R. & Roberts, H. (2018). *Network Propaganda: Manipulation, Disinformation, and Radicalization in American Politics.* Oxford University Press.

Boudry, M, Paglieri, F., & Pigliucci, M. (2015). "The Fake, the Flimsy, and the Fallacious: Demarcating Arguments in Real Life," *Argumentation* 29 (4):431–456.

Bulger, M., & Davison, P. (2018). "The Promises, Challenges, and Futures of Media Literacy," *Journal of Media Literacy Education*, 10(1), 1-21.

Burr, C., Cristianini, N. & Ladyman, J. (2018). "An Analysis of the Interaction Between Intelligent Software Agents and Human Users," *Minds and Machines* 28(4):735-774.

Cassam, Q. (2015). "Stealthy Vices", *Social Epistemology Review and Reply Collective*, 4(10): 19-25.

Cassam, Q. (2016). "Vice Epistemology", *The Monist*, 99:159-80.

Dentith, M. R. X. (2016). "The Problem of Fake News". *Public Reason.* 8 (1-2):65-79.

Douven, I. (2011). "Abduction", *The Stanford Encyclopedia of Philosophy* (Spring 2011 Edition), ed. E. N. Zalta, URL = <http://plato.stanford.edu/archives/spr2011/entries/abduction/>

Flick, D. (2017). "Combatting Fake News: Alternatives to Limiting Social Media Misinformation and Rehabilitating Quality Journalism", 20 SMU Sci. & Tech. L. Rev.

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

Gottfried, J. &. Shearer, E. (2016). "News Use across Social Media Platforms 2016" *Pew Research Center*. http://www.journalism.org/2016/05/26/news-use-acrosssocial-mediaplatforms-2016.

Habgood-Coote, J. (2018). "Stop Talking about Fake News!" *Inquiry: An Interdisciplinary Journal of Philosophy*, 1-33.

Heersmink, R. (2017). "A Virtue Epistemology of the Internet: Search Engines, Intellectual Virtues, and Education," *Social Epistemology*. DOI: 10.1080/02691728.2017.1383530.

Huemer, M. (2009). "Explanationist Aid for the Theory of Inductive Logic," *Brit. J. Phil. Sci.*, 60:345–75.

Jones-Jang, S. M., Mortensen, T., & Liu, J. (2019). "Does media literacy help identification of fake news? Information literacy helps, but other literacies don't," *American Behavioral Scientist*, pp. 1-18. doi:10.1177/0002764219869406.

Larson, H. J. (2018). "The biggest pandemic risk? Viral misinformation." *Nature*, 562(7727):309–309.

Lemish, D., & Lemish, P. (1997). "A much-debated consensus: Media literacy in Israel." In R. Kubey (Ed.), *Media literacy in the information age*. Routledge, 213-228.

Levy, Neil. (2019). "Is Conspiracy Theorising Irrational?" Social Epistemology Review and Reply Collective 8(10):65-76.

Levy, N. (2017). "The Bad News About Fake News." *Social Epistemology Review and Reply Collective* 6(8):20-36.

Lipton, P. (2004). *Inference to the Best Explanation*, 2nd ed. New York: Routledge.

Mason, L.E., Krutka, D.G., & Stoddard, J. (2018). "Media literacy, democracy, and the challenge of fake news*," Journal of Media Literacy Education* 10(2): 1-10.

McIntyre, L. (2018). *Post-Truth*. Cambridge, MA: MIT Press.

Mercier, H. (2018). *Not Born Yesterday*. Princeton University Press.

Nguyen, C.T. (2018)."Epistemic Bubbles and Echo Chambers." *Episteme*. DOI:10.1017/epi.2018.32

Nussbaum, M. (1998). "Whether From Reason Or Prejudice": Taking Money For Bodily Services. *The Journal of Legal Studies, 27*(S2):693-723. doi:10.1086/468040

O'Connor, C & Weatherall, J.O. (2019). *The Misinformation Age*. New Haven, CT: Yale University Press.

Pallavicini, J, Hallsson, B. & Kappel, K. (2018). "Polarization in groups of Bayesian agents," *Synthese*:1-55.

Psillos, S. (2002). "Simply the Best: A Case for Abduction," in A.C. Kakas and F. Sadri (eds.), *Computational Logic: Logic Programming and Beyond*. Berlin: Springer-Verlag, 605-26.

Rini, R. (2017). "Fake News and Partisan Epistemology," *Kennedy Institute of Ethics Journal*, 27 (S2):43-64.

Saez-Trumper, D., Castillo, C. and Lalmas, M. (2013)."Social Media News Communities: Gatekeeping, Coverage, and Statement Bias." In *Proceedings of the 22nd ACM International Conference on Information & Knowledge Management*. https://dl.acm.org/citation.cfm?id=2505623.

Satz, D. (2010). *Why Some Things Should Not Be for Sale: The Moral Limits of Markets*. New York: Oxford University Press.

Silverman, C. (2016). "This Analysis Shows How Fake Election News Stories Outperformed Real News on Facebook." *BuzzFeed News*.

Sober, E. (1994). *From a Biological Point of View: Essays in Evolutionary Philosophy*. New York: Cambridge.

Sober, E. (2002). "Bayesianism—Its Scope and Limits" in R. Swinburne, (ed.), *Bayes' Theorem*, *Proceedings of the British Academy Press*, 113:21–38.

Sosa, E. (1991). *Knowledge in Perspective: Selected Essays in Epistemology*. Cambridge, UK: Cambridge University Press.

Sullivan, E., Sondag, M., Rutter, I., Meulemans, W., Cunningham, S., Speckmann, B., Alfano, M. (2020a). "Can Real Social Epistemic Networks Deliver the Wisdom of Crowds?" In T. Lombrozo, J. Knobe, & S. Nichols (eds.), *Oxford Studies in Experimental Philosophy, Volume 1*. Oxford: Oxford University Press.

Sullivan, E., Sondag, M., Rutter, I., Meulemans, W., Cunningham, S., Speckmann, B., & Alfano, M. (2020b). "Vulnerability in Social Epistemic Networks," *International Journal of Philosophical Studies*:1-23. https://doi.org/10.1080/09672559.2020.1782562

Tanesini, A. (2016). "Teaching Virtue: Changing Attitudes," *Logos and Episteme* 7(4):503-527.

Watson, Lani (2018). "Systematic Epistemic Rights Violations in the Media: A Brexit Case Study," *Social Epistemology* 32(2):88-102.

Zagzebski, L. (1996). *Virtues of the Mind: An Inquiry Into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge, UK: Cambridge University Press.