

Epistemic obligations and free speech

Boyd Millar

Washington University in St. Louis, St.
Louis, Missouri, USA

Correspondence

Boyd Millar, Washington University in St.
Louis, St. Louis, MO, USA.

Email: millar.boyd@gmail.com

Abstract

Philosophical discussions of free speech often focus on moral considerations such as the harm that certain forms of expression might cause. However, in addition to our moral obligations, we also have a distinct set of *epistemic* obligations—and even when a false belief doesn't harm anyone, it constitutes an epistemically bad outcome. Moreover, the existing psychological evidence suggests that human beings are vulnerable to the influence of a wide variety of false claims via a wide variety of psychological mechanisms. Taken together, these facts suggest that there is a purely epistemic justification for restricting the distribution of misinformation: Because each of us has an individual epistemic obligation to avoid unnecessary exposure to misinformation, and because avoiding such exposure is simply too difficult when acting alone, we all have a shared epistemic obligation to establish laws or regulations restricting the widespread distribution of misinformation.

1 | INTRODUCTION

In *On Liberty*, Mill states that we should never restrict public expressions of some claim simply because we're convinced that it's false. Imposing such restrictions, he maintains, would almost certainly leave us worse off. On the one hand, we might be wrong; in which case, by restricting public expressions of the claim at issue we will have deprived ourselves of the “opportunity of exchanging error for truth.” On the other hand, we might be right; but, in that case, we will have deprived ourselves of a different sort of opportunity—the opportunity of achieving the “clearer perception and livelier impression of truth” that is “produced by its collision with error” (Mill, 1859/1977, 229).

Most everything we've learned about human psychology in recent decades suggests that Mill's view is overly optimistic. It turns out that human beings are typically negatively impacted by the falsehoods we encounter: we are liable to believe such falsehoods via a wide variety of psychological mechanisms; and even when we avoid forming false beliefs, most every falsehood we encounter

weakens our grip on the truth.¹ And this widespread vulnerability to false information raises epistemological issues that philosophical discussions of free speech rarely address. Philosophical discussions of free speech are often restricted to a particular moral issue: the potential harms caused by certain forms of expression, such as hate speech or pornography.² But, even when a false (or less accurate) belief doesn't end up harming anyone, it constitutes an epistemically bad outcome. Accordingly, the fact that we are much more vulnerable to false information than Mill realizes suggests that we should consider a distinct approach to the topic: perhaps there is a purely epistemic justification for restricting the public expression of certain falsehoods.³

According to a standard view, in addition to our *moral* and *prudential* obligations, human beings have a distinct set of *epistemic* obligations.⁴ Some of these obligations will be obligations to form or maintain only certain sorts of beliefs (more generally, doxastic attitudes). For instance, perhaps we have an epistemic obligation to believe only propositions that are supported by our evidence. In addition, many of the actions we perform or omit will indirectly influence whether we form or maintain epistemically successful beliefs and avoid forming or maintaining epistemically unsuccessful beliefs (for present purposes, we can understand epistemic success to be either true belief or knowledge). Plausibly, then, some of our epistemic obligations are obligations to perform or omit certain belief-influencing actions.⁵ My argument, stated briefly, is that some of our epistemic obligations to perform belief-influencing actions entail that we ought to impose certain restrictions on the freedom of expression. In particular, each of us has an individual epistemic obligation to avoid unnecessary exposure to certain sorts of falsehoods; and, consequently, those of us who are citizens of democratic states have a shared epistemic obligation to work together to restrict the public expression of such falsehoods.

More specifically, the argument that follows is concerned with *misinformation*: false claims that can be demonstrated to be false on the basis of publicly available evidence. (Rather than attempt to clarify each of the potentially ambiguous terms that this definition employs, I will rely on a supplementary reference-fixing description: false claims of the sort that elicit widespread agreement amongst professional fact-checking organizations.)⁶ I maintain that each of us has an individual epistemic obligation to avoid unnecessary exposure to misinformation. Accordingly, given the prevalence of misinformation distributed via popular internet technologies—specifically, search engines and social media platforms—each of us has an individual epistemic obligation to avoid using such technologies. However, it's simply too difficult for each of us to fulfill this obligation acting alone; and, as such, we all have a shared epistemic obligation to participate in joint actions aimed at restricting the spread of misinformation via the internet technologies at issue.

¹The evidence for these claims about human psychology is outlined below (§3.1).

²Even philosophers who are specifically concerned with misinformation as such focus exclusively on moral considerations and are principally concerned with the harms caused by widespread misinformation: see Mathiesen (2019) and O'Connor and Weatherall (2019, 182–184).

³Levy (2019) defends an argument along these lines. However, Levy's conclusion is much less expansive than the conclusion defended in the present paper, and his argument appeals to very different epistemic considerations.

⁴I won't attempt to argue for this assumption here. For defenses of the view, see, for example, Alston (1985), Feldman (2002), Nottelmann (2007, §1.3), Peels (2017, 101–102), and Brown (2020).

⁵I won't attempt to argue for this assumption either. For defenses of the view, see, for example, Kornblith (1983), Alston (1985), Leon (2002), Nottelmann (2007, chap. 12), Peels (2017, chap. 3), and Lackey (2021).

⁶For discussion of some of the different ways that “misinformation” is standardly defined, see Fallis (2016).

2 | INDIVIDUAL AND SHARED OBLIGATIONS

When the members of some group each possess an individual obligation to contribute to the performance or omission of some joint action, they possess a *shared* or *joint* obligation. For example, some of our individual moral obligations are obligations to perform (or omit) actions that we can perform in conjunction with other people. Accordingly, the individual members of some groups might, for instance, each be morally obligated to work together to help some person in distress; or they might each be morally obligated to refrain from working together to harm some innocent person. So, as long as some of our epistemic obligations are obligations to perform actions that can be performed jointly with other individuals, it follows that we sometimes have shared epistemic obligations. And while some of the actions we are epistemically obligated to perform can't be performed in conjunction with others—for instance, searching our memory or focusing our attention—other actions we are epistemically obligated to perform are the sort that we can (and do) perform in conjunction with others—for instance, searching for information or evaluating evidence.⁷

(Whether moral or epistemic, shared obligations should be distinguished from *collective* obligations. When some group of individuals has a shared obligation to perform or omit some joint action, each group member ought to contribute to that joint action, and each is blameworthy if she fails to do so (though each member may not be equally blameworthy). In contrast, a collective obligation is an obligation that some group of individuals has as a collective entity; as such, that some group has a collective obligation does not entail that any individual group member has the obligation, and so does not entail that any individual member is blameworthy should that obligation be violated. It is controversial whether any collective entities have obligations above and beyond those of their members; conversely, the suggestion that the individual members of a group sometimes share an obligation to contribute to some joint action is not controversial.)⁸

However, simply acknowledging that groups of individuals sometimes possess shared epistemic obligations will not suffice for present purposes; what's required is a rule for determining *when* groups possess shared obligations. We don't need to establish anything like necessary and sufficient conditions for possessing a shared obligation; but, the following argument requires that individual obligations entail shared obligations under certain specific conditions. Specifically, the argument relies on what I'll call the *Shared Obligation Principle*:

If some action (omission) is an obligation for the individual members of some group, but fulfilling that obligation would be so difficult for each individual acting alone that most cannot be reasonably expected to do so, and contributing to some corresponding joint action (omission) would be sufficiently easy that most individual members can be reasonably expected to do so, then these individuals have a shared obligation to perform (omit) that joint action.

We should note two additional details. First, since there will often be a number of different ways in which (or methods by which) the obligatory joint action can be performed, the principle implies only that group members have a shared obligation to perform the joint action in one of the available ways. Second, I claim only that the principle applies in cases where the corresponding joint action would not itself constitute a violation of some obligation of equal or greater importance.

⁷For further discussion of shared epistemic obligations, see Millar (2021).

⁸For discussion, see, for example, Feinberg (1968), Sverdlik (1987), May (1992), Narveson (2002) and Miller and Makela (2005).

The Shared Obligation Principle is supported by some plausible assumptions concerning the relation between individual and shared obligations. When an individual is obligated to perform a given action and she is able to perform that action either by acting alone or by acting in conjunction with other people, we can't assume that she is obligated to contribute to some joint action; but, if in such a case, the methods by which she might perform this obligatory action while acting alone are all eliminated, then she is obligated to contribute to some joint action. This claim is an instance of a more general principle: when all other available methods of fulfilling some obligation are eliminated, whatever method remains becomes obligatory (at least if employing that method does not itself violate some obligation). For instance, suppose Ben ought to provide someone a meal and that, under the circumstances, there are exactly two methods available to him: use the money in his wallet to purchase a meal, or hand over his sandwich. If Ben's wallet is now stolen, clearly he ought to hand over his sandwich. Next, suppose Ben and his two brothers each ought to provide someone a meal; and suppose that each has enough money in his wallet to purchase a meal, and that each is holding a third of a sandwich. If Ben and his brothers now each has his wallet stolen, clearly each ought to cooperate with his brothers in order to provide a meal by combining their partial sandwiches.⁹ In such a case, if Ben fails to fulfill his obligation, he is still fully blameworthy even though it is impossible to fulfill his obligation while acting alone—and the only explanation for why he is fully blameworthy is that he is obligated to contribute to the joint action. So, the fact that it has become impossible for each sibling to fulfill his obligation while acting alone entails that each ought to contribute to the corresponding joint action. That is, when all individual methods are eliminated, the available joint methods become obligatory.

Joint methods also become obligatory when individual methods become extremely difficult, though not strictly impossible. In particular, in the sort of cases at issue, contributing to some joint action becomes obligatory when fulfilling the obligation while acting alone is so difficult that the individual cannot be reasonably expected to do so—in other words, in cases where the difficulty fulfilling the obligation individually would provide an excuse that at least significantly mitigates blame.¹⁰ For instance, suppose you encounter a child who is upset at having lost a small object of significant sentimental value; and suppose that the only way to locate this object is to carefully search an enormous field. In this case, it's plausible that you have an obligation to help the child find the object. But fulfilling this obligation while acting alone would be so difficult that you could not be reasonably expected to do so—it will take hours to search the field by yourself, and you have other important things to do. If, then, searching the field by yourself were the only method by which you could fulfill your obligation, you would have an excuse that at least significantly mitigates blame should you fail to do so. But, next, suppose that there are dozens of other people nearby, each of whom is in the same situation as yourself; and suppose that if each of you were to focus on a small area of the field, the search would be completed rather quickly. Clearly, you and these other individuals have a shared obligation to perform the relevant joint action. In such a case, the difficulty fulfilling the obligation when acting alone no longer provides an excuse should you refuse to help the child find the object. Rather, you would be

⁹Following Hindriks (2019), one might claim that, initially, each member of this group is only obligated to encourage his brothers to commit to the relevant joint action; and that, each member is only obligated to contribute to the joint action once his brothers have committed themselves. However, this account of the case is problematic for at least two reasons. First, it entails that the group members are obligated to pursue the means of achieving some outcome that they are not obligated to achieve. Second, it is unnecessarily complicated. The simplest thing to say about the present case is that Ben ought to provide the meal, but that the means required to fulfill this obligation are complex (as is very often the case with purely individual obligations). That is, fulfilling his obligation will require that he perform a complex action—the first step of this action will involve talking to his brothers, while the second will involve handing over his partial sandwich.

¹⁰For discussion of such excuses, see, for example, Franklin (2013, §3), Peels (2017, 134–137), and Sliwa (2019, 42–43).

fully blameworthy for failing to fulfill your obligation—and the only explanation for why you are fully blameworthy is that you ought to contribute to the joint action (and have no excuse for failing to do so).

3 | THE EPISTEMIC OBLIGATION TO AVOID MISINFORMATION

Simply acknowledging that some of our epistemic obligations are obligations to perform or omit belief-influencing actions will not suffice for present purposes; rather, the following argument relies on the assumption that each of us has an epistemic obligation to engage in a certain sort of belief-influencing behaviour. Specifically, the argument relies on the assumption that each of us has an epistemic obligation to avoid unnecessary exposure to misinformation—especially misinformation that we would encounter repeatedly and misinformation consistent with what we already believe. (By “unnecessary” I mean something like “except in special circumstances.” In special circumstances, the only way to acquire certain information, or to develop some particular cognitive skill, might involve exposing oneself to misinformation. Doing so in such circumstances would not violate the proposed epistemic obligation.)

The argument for the claim that each of us has an epistemic obligation to avoid unnecessary exposure to misinformation proceeds in two stages. First, I review the empirical evidence suggesting that encountering false information negatively impacts our doxastic attitudes in systematic ways—especially when we encounter falsehoods repeatedly and when those falsehoods are consistent with what we already believe. Second, I argue that these psychological facts entail that each of us possesses the epistemic obligation at issue.

3.1 | Our vulnerability to false information

There is considerable empirical evidence for the following claims. First, the falsehoods we encounter are likely to influence our doxastic attitudes in problematic ways—especially when we encounter a given falsehood repeatedly or when a given falsehood is consistent with what we already believe. Second, there is nothing we can do as individuals to eliminate our vulnerability to false information. And, third, once our doxastic attitudes have been influenced by some falsehood, it is unlikely that this influence can be largely reversed or corrected.

Whenever we encounter some claim and it isn't immediately clear whether that claim is true or false, each of us exhibits a *truth bias*—a tendency to accept such claims as true.¹¹ This fact should be evident from reflection on your own experience: rejecting something someone says or something you read is much less common than passive acceptance. But this tendency is more pervasive than reflection might suggest—it is sufficiently strong that we often accept claims of unknown truth value even when we have very good reasons not to do so. For instance, Pantazi et al. (2018) discovered striking evidence of this phenomenon when they presented subjects with a series of statements regarding a fictional crime, and then asked them to rate the perpetrator's dangerousness and to recommend a prison sentence. Some of the presented statements were explicitly labelled as true, while others were

¹¹For a review of the relevant evidence, see Brashier and Marsh (2020, 501–502). Two important points to note: the claim regarding the truth bias endorsed here does not entail that comprehending a statement involves accepting it (see, for example, Gilbert et al., 1993 and Mandelbaum, 2014); and it is consistent with the fact that individuals typically avoid believing claims that are either incompatible with stored knowledge (see Richter et al., 2009) or that are both informative when false and explicitly labelled as false (see Hasson et al., 2005).

explicitly labelled as dubious—dubious in the sense that subjects were informed that these statements were taken from entirely unrelated crime reports.¹² Nonetheless, the results were that subjects had a tendency to misremember dubious statements as true (but not vice versa), and to judge the perpetrator more severely when the dubious statements they had encountered increased the severity of the crime. Crucially, participants demonstrated a significant tendency to treat the dubious claims as true even when not subjected to distractions of any kind.

Thanks to our natural truth bias, then, the beliefs we form on the basis of the falsehoods we encounter are likely to be false—at least so long as those falsehoods are not flatly inconsistent with our existing beliefs. Moreover, we are especially likely to accept the claims we encounter as true when those claims are consistent with what we already believe or when we encounter those claims repeatedly. Each of us checks the claims we encounter against what we already believe in a spontaneous, automatic fashion.¹³ Such a procedure is perfectly rational; but, the result is that we're likely to believe falsehoods that happen to be consistent with our existing beliefs—for instance, individuals are much more likely to believe even implausible fake news headlines that are consistent with, rather than inconsistent with, their political beliefs.¹⁴ Somewhat more troublingly, thanks to our *confirmation bias*, each of us has a tendency to accept claims consistent with our beliefs relatively uncritically.¹⁵ For instance, individuals accept evidence and arguments of a certain quality much more readily when that evidence supports rather than challenges their beliefs.¹⁶ In addition, each of us is more likely to believe a given claim the more frequently we encounter it—a phenomenon known as the *truth effect*.¹⁷ And because each separate time we encounter some claim increases our confidence that that claim is true, sheer repetition can cause us to believe things we shouldn't.¹⁸ For instance, individuals judge familiar but implausible fake news stories to be more accurate than unfamiliar real news stories.¹⁹ And Henkel and Mattson (2011) showed that individuals who read a particular statement on three separate occasions over the course of two weeks, and who were given explicit advance warning that the source of the information was unreliable, still judged that statement to be either probably or certainly true almost 70% of the time.

Importantly, in cases where you encounter some falsehood and manage to avoid believing it (perhaps because it is inconsistent with your existing beliefs), this encounter will typically still negatively influence your doxastic attitudes. Research on the truth effect has shown that your grip on the truth is weakened each time you encounter some falsehood inconsistent with what you know to be true. For instance, researchers have found that even when some particular individual knows that the Pacific Ocean is the largest ocean on Earth, each time she reads the statement that “the Atlantic Ocean is largest ocean on Earth,” her rating of how inaccurate that statement is decreases.²⁰ Similarly, individuals who recognize that a particular implausible fake news story is inaccurate rate that story to

¹²When Pantazi et al. (2018) describe their experiment, they claim that some statements were explicitly labelled as true and the rest explicitly labelled as *false*. However, because the experiment was based on Gilbert et al. (1993), subjects were told that the “false” statements were taken from unrelated crime reports; and, as such, these experiments' instructions made it clear to subjects that the so-called “false” statements could not be assumed to be either true or false.

¹³For a review of the relevant evidence, see Mercier (2017, 104–105).

¹⁴Pennycook et al. (2018) and Pennycook and Rand (2019).

¹⁵For reviews of the relevant evidence, see Nickerson (1998) and Ditto et al. (2019).

¹⁶See, for example, Ditto and Lopez (1992) and Taber and Lodge (2006).

¹⁷For a review of the relevant evidence, see Unkelbach and Koch (2019).

¹⁸DiFonzo et al. (2016).

¹⁹Pennycook et al. (2018).

²⁰Fazio et al. (2015). See also Fazio et al. (2019).

be significantly less inaccurate each time they encounter it.²¹ In addition, simply encountering false claims—even false claims you don't believe—can undercut the positive effects that encountering true claims would otherwise have on your doxastic attitudes. For instance, van der Linden et al. (2017) showed that reading the statement that “97% of climate scientists have concluded that human-caused climate change is happening” will significantly improve the accuracy of an individual's estimate of the current level of scientific consensus regarding climate change; unless, that is, that individual subsequently reads the false statement that “there is no consensus on human-caused climate change”—in which case, encountering this falsehood completely eliminates the benefits of having first read the true statement. Moreover, encountering claims that you recognize to be false frequently exerts a problematic influence on your reasoning and judgement. For instance, research has shown that individuals often reproduce the false claims they have read in stories when answering trivia questions, even when those falsehoods flatly contradict facts they know to be true.²²

While there are steps we can take to partially mitigate this negative influence of false information, the falsehoods we encounter are likely to have a significant negative influence on our doxastic attitudes no matter what we do. As we've already seen, you're likely to be influenced by false claims even when those claims are incompatible with what you know, and even when they are explicitly labelled as dubious. You might think that you could reduce or even eliminate this influence by taking appropriate precautions—perhaps by reading slowly and carefully. However, research has shown that reading slowly and carefully does not reduce the consequences of encountering false information.²³ Or, you might think that you could protect yourself against the influence of false information by developing certain cognitive skills, or by adopting a sceptical mindset. However, research has shown that individual differences in cognitive ability and personality have little or no bearing on one's susceptibility to confirmation bias or the truth effect.²⁴ Individuals with a greater propensity towards “analytic” rather than “intuitive” thinking are somewhat better able to identify certain kinds of falsehoods—specifically, implausible fabricated news headlines designed to get attention online.²⁵ So, cultivating your propensity towards analytic thinking might make you somewhat better able to spot fake news headlines specifically; but such a strategy would still leave you liable to believe false information consistent with your existing beliefs,²⁶ would provide no protection against the influence of repetition,²⁷ and wouldn't eliminate the other negative impacts that encounters with false information have on your doxastic attitudes (reviewed above). Alternatively, you might think that it would help to acknowledge and attempt to actively suppress your cognitive biases; but the available evidence suggests that such attempts would produce only minor improvements.²⁸ For instance, Nadarevic and Aßfalg (2017) provided individuals with a description of the truth effect, told them that half of the statements they were about to read were false, and explicitly instructed them not to allow repetition to influence their judgements regarding the

²¹Pennycook et al. (2018).

²²Fazio et al. (2013).

²³See, for example, Fazio and Marsh (2008), Fazio et al. (2013), and Rapp (2016).

²⁴See, for example, DiFonzo et al. (2016), Mercier (2017), and De keersmaecker et al. (2020).

²⁵Bronstein et al. (2019) and Pennycook and Rand (2019, 2020).

²⁶For instance, Pennycook and Rand found that “there was an overall tendency for Democrats/Clinton supporters to rate Democrat-consistent fake news as more accurate than Republican consistent fake news (and vice versa for Republicans/Trump supporters)” (2019, 47).

²⁷Pennycook and Rand (2020).

²⁸For discussion of why personal debiasing strategies are very likely to be unsuccessful, see Levy (2012, 597–598) and Ahlstrom-Vij (2013, §1.4).

truth of these statements—yet even under these ideal conditions, the resulting judgements exhibited a significant truth effect.²⁹

One final method you might think you could employ to protect yourself against the influence of false information would be to seek out accurate information—perhaps by utilizing multiple reliable sources, or actively researching dubious claims. However, research has shown that once our doxastic attitudes have been influenced by some falsehood, this influence cannot be largely reversed or corrected.³⁰ In instances where you believe some false information that you've encountered and later encounter a correction, if the correction is directly incompatible with your deeply held beliefs, typically you will simply reject the correction and maintain your false belief.³¹ Of course, most of the corrections you encounter will not fall into this category. Yet, in most cases in which you encounter false information and later encounter a correction, the accurate information will only partially counteract the influence the false information has had on your doxastic states—and, even then, this partial improvement will be only temporary. Research on the effectiveness of reading corrections of false information has shown that reading corrections typically improves the accuracy of an individual's doxastic attitudes; but, the improvement is only partial: individuals who read the correction tend to have more accurate doxastic attitudes than they would have if they hadn't read the correction, but not so accurate as they would have if they had never encountered the false information in the first place.³² Moreover, research has also shown that this beneficial effect of reading corrections tends to diminish rather quickly. For instance, Porter and Wood (2019, 42–43) found that, in a wide range of cases, the beneficial effects of reading corrections of false information “were not statistically distinguishable from zero after three days.”³³

In addition, even in ideal cases where we encounter a correction and find it fully convincing, the false information that we previously encountered but now reject will continue to influence our doxastic attitudes in systematic ways—a phenomenon known as the *continued influence effect*.³⁴ Recent neuroimaging studies suggest that the source of the problem is that when you reject previously encountered false information on the basis of encountering accurate information, the false information is not simply deleted from the brain—rather, it continues to be stored in memory and continues to be activated when you deliberate on relevant topics.³⁵ So, for instance, Green and Donahue (2011) asked subjects to read a news story touching on some important social issue, after which they revealed to some of these subjects that the story was intentionally fabricated. After revealing this fact to some subjects, they then asked a series of questions to gauge the influence reading the story had had on various “story-relevant beliefs” (such as beliefs concerning the effectiveness of certain social programs); they found that the story had the *very same* influence on the story-relevant beliefs of subjects who were told it was intentionally fabricated, as it did on the story-relevant beliefs of subjects who were not told that it was fabricated.³⁶

²⁹The most effective strategies for resisting the influence of false information are impossible to employ in most cases because they require that individuals be able to recognize that the relevant claims are false when they first encounter them. Specifically, individuals can largely avoid being influenced by false statements when they actively correct or explicitly rate the accuracy of those claims as they are reading: see Rapp et al. (2014) and Brashier et al. (2020).

³⁰For an extended defense of this claim, see Levy (2017).

³¹Ecker and Ang (2019). For a review of the evidence concerning the conditions under which beliefs persist in the face of corrections, see Jerit and Zhao (2020, 81–85).

³²See, for example, Chan et al. (2017) and Walter et al. (2020).

³³See, also, Swire et al. (2017).

³⁴Johnson and Seifert (1994). For a review relevant evidence, see Lewandowsky et al. (2012).

³⁵Gordon et al. (2019).

³⁶See also Thorson (2016).

3.2 | The implications of our vulnerability to false information

The foregoing evidence suggests that whenever we encounter false information it is likely to have a significant negative influence on our doxastic attitudes, and the only way for us to avoid this negative influence is to avoid encountering false information altogether. It isn't possible to avoid encountering falsehoods as such—after all, many false claims are supported by the balance of the existing evidence. But misinformation is different in this regard. When there is publicly available evidence that demonstrates that some claim is false, we are in a position to know that that claim is false; and so, misinformation is avoidable by its very nature. Accordingly, our vulnerability to false information entails that each of us has an epistemic obligation to avoid misinformation. More specifically, each of us has an epistemic obligation to avoid unnecessary exposure to misinformation—especially misinformation that we would encounter repeatedly and misinformation consistent with what we already believe. In practice, in order to avoid exposure to misinformation we must avoid relying on sources that regularly dispense misinformation; and so, each of us ought to avoid relying unnecessarily on sources that regularly dispense misinformation—especially when by relying on a given source we would encounter the same misinformation repeatedly or encounter misinformation consistent with what we already believe.³⁷

If you want to establish that a certain action is obligatory, a natural strategy is to show that each of the most plausible theories of right action entails that that action is obligatory. For instance, a simple way to establish that we ought not to kill human beings unnecessarily is to show that the most plausible consequentialist, deontological, and virtue-based accounts of right action each entails that we ought not to do so. While theories of epistemically obligatory action are not as well developed as theories of morally obligatory action, it isn't difficult to determine what epistemological analogues of the standard moral theories will look like. Accordingly, there are good reasons to maintain that the most plausible theories of epistemically obligatory action all entail that each of us has an epistemic obligation to avoid unnecessary exposure to misinformation.

First, one might characterize what makes a belief-influencing action epistemically obligatory in consequentialist terms.³⁸ According to this theory, some action is an epistemic obligation if and only if it leads to epistemically good outcomes. Plausibly, on any reasonable account of epistemically good outcomes—for example, forming true beliefs and avoiding false beliefs, or acquiring knowledge and avoiding ignorance—relying on sources that regularly dispense misinformation will tend to result in bad outcomes. The misinformation you encounter while relying on such sources is likely to result in at least some false beliefs—especially if you encounter the same falsehoods repeatedly, or if these falsehoods are consistent with what you already believe. Misinformation that doesn't directly precipitate false beliefs is likely to weaken your grip on the truth and undermine the beneficial impact that encountering accurate information would otherwise have. Moreover, these negative consequences are not likely to be reversed: if you later encounter a correction of some piece of misinformation, you

³⁷Levy and Mandelbaum (2014) defend a related thesis; but their argument differs from the present argument in a number of important respects. First, they defend a more restricted conclusion: they maintain that individuals have an epistemic obligation only to weigh the potential costs against the potential benefits before consulting a problematic source. Second, their argument presupposes a consequentialist account of our obligations (more on this issue below). Third, their argument relies heavily on Gilbert's view that comprehending a statement involves accepting it (see n11 above).

³⁸The label "epistemic consequentialism" typically refers to views focused on the consequences of holding a certain belief or set of beliefs, rather than the consequences of belief-influencing actions: see, for example Ahlstrom-Vij and Dunn (2018). Nonetheless, arguments that seek to establish that certain belief-influencing actions are epistemic obligations often appeal to the fact that such actions promote epistemically good outcomes: see, for example, Hall and Johnson (1998), Nottelmann (2007, 182–186), Levy and Mandelbaum (2014), Peels (2017, 99–100), and Rettler and Rettler (2021, §4).

might automatically reject it; and if you don't reject it, the correction's beneficial influence is likely to be merely partial and temporary. And, even in cases in which you fully endorse the correction, the misinformation will continue to be stored in memory and influence the beliefs you form on related topics. Of course, a source of the sort at issue needn't dispense misinformation exclusively; so, by relying on such a source you may well acquire many more true beliefs than false beliefs. But, those epistemic benefits can be achieved by relying on alternative sources of information—sources that don't regularly dispense misinformation. As such, except in special circumstances, you are guaranteed to achieve better epistemic outcomes by avoiding sources that regularly dispense misinformation. So, a consequentialist account of epistemically obligatory action—whether focused on individual acts or rules—entails that you ought to avoid relying unnecessarily on sources that regularly dispense misinformation.

Second, one might characterize what makes a belief-influencing action epistemically obligatory in deontological terms. According to this theory, some action is an epistemic obligation if and only if it conforms to our epistemic duties. Plausibly, on any reasonable account of our epistemic duties, relying on sources that regularly dispense misinformation will typically violate those duties. For instance, a natural proposal is that we have an epistemic duty not to cause epistemic harm. Regardless of how one characterizes epistemic harm, relying on sources that regularly dispense misinformation will cause such harm since, as we've seen, doing so undermines the accuracy of our beliefs in systematic ways. An epistemic deontologist might claim that the duty not to cause epistemic harm can be overridden by competing epistemic duties—perhaps the duty to acquire true beliefs, or the duty to develop certain cognitive skills. But given the sheer extent of the harms that exposure to misinformation causes, and given that it will typically be possible to fulfill our positive epistemic duties by relying on alternative sources of information, relying on sources that regularly dispense misinformation will only be consistent with epistemic duty under special circumstances.

Alternatively, an epistemic deontologist might adopt the Kantian view that there is only a single rule governing belief-influencing action. An epistemic analogue of Kant's categorical imperative might be something like: act in such a way that you treat true belief (or accurate belief, or knowledge) always as valuable for its own sake.³⁹ The characteristic feature of a Kantian view is that it prohibits the sorts of epistemic trade-offs that a consequentialist view permits⁴⁰—for instance, the Kantian insists that you ought not to acquire false beliefs in order to acquire other true beliefs. So, if you rely on the best possible sources available to you and those sources happen to include a number of false claims, you have not violated your epistemic obligations because you did everything you could to acquire only true beliefs. But, if you rely on sources that regularly dispense misinformation because those sources are convenient, or engaging, or even because you'll achieve more true beliefs in the process, you are permitting yourself to acquire false (or less accurate) beliefs in order to achieve some other goal—and thereby disregarding the value of true belief. Accordingly, the Kantian account of our epistemic obligations seems to entail that you ought not to rely on sources that regularly dispense misinformation under *any* circumstances. But, in any case, it's clear that any plausible deontological account of epistemically obligatory action entails that you ought to avoid relying *unnecessarily* on sources that regularly dispense misinformation.

Finally, one might characterize what makes a belief-influencing action epistemically obligatory in terms of epistemic virtues. According to this theory, some action is an epistemic obligation if and

³⁹Another way to express the point—drawing on Sylvan (2020)—would be that our actions should always manifest respect for truth or accuracy. Sylvan is specifically concerned with whether beliefs (rather than actions) manifest respect for accuracy; but it's natural to assume that belief-influencing actions can manifest or fail to manifest respect for accuracy as well.

⁴⁰For discussion, see Berker (2013).

only if it is an action that an epistemically virtuous person would characteristically perform. Plausibly, on any reasonable account of the epistemic virtues, an epistemically virtuous person will not rely on sources that regularly dispense misinformation. According to standard accounts, an epistemically virtuous person is one who is motivated by the desire to acquire true beliefs (or knowledge) and avoid false beliefs.⁴¹ Such an individual will often rely on sources of information that happen to include false claims—after all, often the best sources available include such claims. But such an individual will not rely on sources that regularly dispense misinformation—especially when by relying on a given source she would encounter the same misinformation repeatedly, or encounter misinformation consistent with what she already believes. For instance, an epistemically virtuous person is epistemically cautious and possesses epistemic humility.⁴² So, she is aware that her cognitive limitations make her vulnerable to misinformation no matter how carefully she scrutinizes what she reads, how thoroughly she investigates dubious claims, or how hard she tries to suppress her biases; and consequently, she takes that precaution which is required for her to maintain true beliefs and avoid false beliefs—namely, she avoids sources likely to expose her to misinformation. Conversely, to rely on a source that regularly dispenses misinformation because you believe that you aren't likely to be influenced by such falsehoods is to exhibit epistemic overconfidence; and to rely on such sources because you find them convenient or engaging, or because you simply don't care about the risks, is to exhibit epistemic carelessness. Of course, an epistemically virtuous person will rely on a problematic source under special circumstances—but only when she has epistemically creditable grounds for doing so. Accordingly, a virtue-based account of epistemically obligatory action entails that you ought to avoid relying unnecessarily on sources that regularly dispense misinformation.

(A potential objection to the preceding argument is that most people aren't aware of the extent of their vulnerability to false information. However, this objection assumes that individuals must recognize their vulnerability to false information in order to possess the epistemic obligation at issue—and we should reject that assumption. In particular, non-culpable ignorance is more plausibly regarded as providing an excuse for violating this obligation rather than eliminating it. Additionally, we could respond to this objection by appealing to the distinction between *objective* and *subjective* obligations⁴³—that is, we could maintain that individuals who are aware of their vulnerability to false information have both a subjective and an objective obligation to avoid misinformation, while individuals who are ignorant of this vulnerability have only an objective obligation to avoid misinformation. The Shared Obligation Principle seems no less plausible with this distinction in place. Much more substantial amendments would need to be made to the argument presented here in order to accommodate those who insist that our only epistemic obligations to perform belief-influencing actions are subjective obligations.)⁴⁴

4 | THE EPISTEMIC OBLIGATION TO AVOID USING POPULAR INTERNET TECHNOLOGIES

Each of us, then, has an epistemic obligation to avoid relying unnecessarily on sources that regularly dispense misinformation—especially when by relying on a given source we would encounter the same

⁴¹See, for example, Montmarquet (1993) and Zagzebski (1996).

⁴²For discussion, see, for example, Roberts and Jay Wood (2007, chap. 8) and Whitcomb et al. (2017).

⁴³For discussion, see, for example, Olsen (2017).

⁴⁴Hall and Johnson (1998, 130–131), Levy and Mandelbaum (2014, 15), and Peels (2017, 104–105) endorse views along these lines. For further discussion of this issue, see n60 below.

misinformation repeatedly, or encounter misinformation consistent with what we already believe. But certain popular internet technologies—specifically, social media platforms such as Facebook, Twitter, and YouTube, and search engines such as Google and Bing—dispense great quantities of misinformation. Moreover, these social media platforms and search engines employ personalization algorithms to filter and organize the information that users encounter. As a result, when you rely on such sources, not only are you likely to encounter misinformation, but you are likely to encounter the same misinformation repeatedly, and that misinformation is likely to be consistent with what you already believe.

Research has demonstrated that a significant portion of the information dispensed via popular social media platforms is misinformation. For instance, a survey of all news stories posted on Twitter between 2006 and 2017 found that news stories presenting misinformation typically “reached far more people” than did true news stories.⁴⁵ Misinformation is similarly prevalent on Facebook: for instance, a study focused on 100 specific items of potentially harmful COVID-19 misinformation circulating on Facebook found that, over a roughly two-month period in early 2020, posts including this misinformation were viewed an estimated 117 million times, and videos including this misinformation were watched 47 million times.⁴⁶ YouTube is no better: for instance, a recent study found that amongst the top 100 global-warming-related videos on YouTube, those presenting misinformation accounted for 20% of all views.⁴⁷ In addition, social media includes an extra layer of misinformation in the form of social media user comments. Anspach and Carlson (2020) found that even when Facebook and Twitter users shared accurate news stories, the comments they added regularly misrepresented those stories; and, as a result, individuals exposed to such comments were more misinformed than individuals who were not exposed.

Research concerning the distribution of misinformation via search engines has produced similar results. A recent study demonstrated that Google and Bing’s top 50 search results for a variety of queries included many sources of misinformation.⁴⁸ Bing proved to be particularly problematic—roughly 20% of the sources included in its top 50 results for the relevant queries presented misinformation.⁴⁹ But while Google may be less problematic, it nonetheless regularly dispenses misinformation. For instance, in December 2016, Google’s first result for the query “did the holocaust happen” was a neo-Nazi webpage listing the “top 10 reasons the holocaust didn’t happen.”⁵⁰ A few days following the 2016 Presidential election, Google’s first result for the query “final election results” was an article claiming that Donald Trump had won the popular vote.⁵¹ And around the same time, in response to the query “is Obama planning a coup,” Google explained in a special box at the top of the page (a Google *featured snippet*) that, yes, Obama was planning to stage a coup at the end of his Presidential term.⁵² Another way that Google prominently features misinformation is via its autocomplete function. For instance, someone typing “climate change is” into Google will typically be presented with suggestions such as “climate change is not real” and “climate change is a hoax.”⁵³

The quantity of misinformation dispensed via social media platforms and search engines is particularly problematic given that the information you encounter when using such sources is shaped

⁴⁵Vosoughi et al. (2018, 1149). For further evidence of the significant quantity of misinformation spread via Twitter, see Grinberg et al. (2019) and Bovet and Makse (2019).

⁴⁶Avaaz (2020a).

⁴⁷Avaaz (2020b).

⁴⁸Bush and Zaheer (2019).

⁴⁹Ibid.

⁵⁰Bhattacharya (2016).

⁵¹Bump (2016).

⁵²Ohlheiser (2017).

⁵³Lapowsky (2018). For discussion, see Miller and Record (2017, §§3–4).

by what you already believe. When using social media, the information you're presented with is determined primarily by who your friends are and whom you follow. And since the individuals who make up your social network are likely to share your beliefs—for example, liberals tend to have more liberal Facebook friends, and conservatives tend to have more conservative friends—the information you encounter is likely to be consistent with your beliefs.⁵⁴ In addition, search engines and social media platforms use personalization algorithms to filter and organize the information you're presented with so that you're more likely to encounter information you find appealing.⁵⁵ Research has shown that individuals using search engines and social media to obtain information spend significantly more time reading attitude-consistent than attitude-inconsistent articles.⁵⁶ Perhaps unsurprisingly, this personalized filtering of information has been shown to play a primary role in the distribution of misinformation on social media: a particular false claim spreads widely on social media because it's able to reach precisely those individuals most likely to find it plausible.⁵⁷

Ultimately, then, because each of us has an epistemic obligation to avoid relying unnecessarily on sources that regularly dispense misinformation—especially when by relying on a given source we would encounter the same misinformation repeatedly, or encounter misinformation consistent with what we already believe—each of us has an epistemic obligation to avoid obtaining information via social media and search engines unnecessarily. First, social media platforms and search engines regularly dispense misinformation. And, second, because they combine large quantities of misinformation with the personalized filtering of information, when we use these popular internet technologies we are likely to encounter the same misinformation repeatedly, and we are likely to encounter misinformation consistent with what we already believe.

Moreover, relying on social media platforms and search engines is rarely necessary because we can almost always obtain the information we're after via some other source. We can acquire information regarding news and current events from newspapers and magazines; we can keep up to date with the lives of our friends and family by communicating with them directly; and we can conduct any necessary research by consulting some trustworthy online resource or visiting a library. Of course, special circumstances will arise in which it is impossible for some individual to acquire some information except by relying on search engines or social media; and an individual who relies on these sources under such circumstances does not violate the epistemic obligation at issue. But, for most of us, it will be epistemically permissible to acquire information via social media or search engines only rarely.⁵⁸

5 | THE SHARED EPISTEMIC OBLIGATION TO RESTRICT MISINFORMATION

According to the Shared Obligation Principle, when the individual members of some group possess an obligation that is so difficult to fulfill when acting alone that most cannot be reasonably expect to do so, they ought to perform some corresponding joint action—at least so long as contributing to that

⁵⁴See, for example, Gaines and Mondak (2009), Himelboim et al. (2013), and Halberstam and Knight (2016).

⁵⁵Bakshy et al. (2015). For discussion, see Miller and Record (2013).

⁵⁶See Knobloch-Westerwick et al. (2015) and Dylko et al. (2017).

⁵⁷See, for example, Del Vicario et al. (2016) and Shin et al. (2017).

⁵⁸In principle, it is possible to use at least certain social media platforms in such a way that one is highly unlikely to encounter any misinformation. However, social media used in such a restrictive manner would lose much of its purpose—for instance, one could no longer use social media to maintain relationships with friends and family. Very specific kinds of search engine use would also presumably be safe; if so, then there is nothing epistemically problematic about using search engines in the relevant way.

joint action is sufficiently easy for the relevant individuals. We've now seen that each of us has an epistemic obligation to avoid relying unnecessarily on sources that regularly dispense misinformation. Given the way that social media platforms and search engines currently function, this fact entails that each of us has an epistemic obligation to avoid obtaining information via social media and search engines unnecessarily. But, in practice, each of us will find fulfilling this obligation rather difficult; in fact, fulfilling this obligation while acting alone is sufficiently difficult that most of us can't be reasonably expected to do so.

It typically isn't necessary for us to rely on social media and search engines because we have more reliable methods of obtaining information; however, these alternative methods require significantly more time, money, and effort—additional time, money, and effort that makes it difficult for most individuals to avoid relying on the internet technologies at issue. While you can acquire information regarding news and current events from newspapers and magazines, subscribing to such publications costs substantially more money than relying on Google and Facebook does; while you can keep up to date with the lives of your friends and family by communicating with them directly, doing so requires substantially more time and effort than utilizing social media does (and because many of your friends and family will be less inclined to communicate with you via alternative means, many of your relationships are likely to suffer); and while you can conduct any necessary research by consulting some trustworthy online resource or visiting a library, doing so typically requires substantially more time and money than relying on Google does. Accordingly, most people can't be reasonably expected to expend the additional time, money, and effort required to avoid relying on social media and search engines—especially over extended periods of time.

It's important to keep in mind that what can be reasonably expected of someone under certain circumstances depends on the stakes. For example, while you can't be reasonably expected to walk several miles to spare someone a mild inconvenience, you can be reasonably expected to do so to save someone's life. In the case at hand, the epistemic stakes are relatively low: it's not as though by obtaining information via social media and search engines you will suffer an epistemic catastrophe, such as forming predominantly inaccurate beliefs. Yet, the long-term costs associated with permanently avoiding social media and search engines are substantial; and so, it would be unreasonable to expect most people to expend the additional time, money, and effort required to do so given the epistemic stakes. To put the point another way: the difficulty that most people would have fulfilling the epistemic obligation at issue provides an excuse that at least significantly mitigates blame. If you rely on social media and search engines due the additional time, money, and effort required to acquire information via other means—perhaps because, like most people, you have to manage your time and money carefully—you deserve significantly less blame for violating your epistemic obligation than you would if fulfilling it wasn't particularly difficult.

In contrast, those of us who are citizens of democratic states have methods of fulfilling our epistemic obligation to avoid misinformation that require very little time, money or effort—specifically, we can work with our fellow citizens to fulfill this obligation via some joint action. Crucially, we only have an epistemic obligation to avoid relying on social media and search engines so long as they continue to regularly dispense misinformation. As such, refusing to rely on such sources is just one way to fulfill our obligation to avoid unnecessary exposure to misinformation; alternatively, we could simply stop them from dispensing so much misinformation. None of us has the ability to stop social media platforms and search engines from regularly dispensing misinformation while acting alone; but the citizens of democratic states have precisely this ability in virtue of their influence over public policy. In general terms, the joint action required to stop the relevant entities from dispensing misinformation consists in the citizens of democratic states, via our political representatives, enacting appropriate laws or regulations. The contribution that each individual citizen must make to this

process doesn't require much in the way of time, money, or effort: each of us need only write to our elected representatives, sign any relevant petitions, and vote for politicians who support the relevant policies. So, the contribution that each of us would need to make to the requisite joint action is sufficiently easy that each of us can be reasonably expected to do so—for instance, the time, money, or effort involved in signing a petition, or voting, would not significantly mitigate the blame any one of us would deserve should we fail to do our part.⁵⁹ Consequently, the Shared Obligation Principle entails that we have a shared epistemic obligation to perform some joint action that would prevent social media platforms and search engines from distributing misinformation.⁶⁰

There are a number of different methods by which social media platforms and search engines might be prevented from distributing misinformation; accordingly, the Shared Obligation Principle entails only that we have a shared epistemic obligation to employ one such method. That is, the Shared Obligation Principle entails that we ought to adopt laws or regulations the result of which will be that a typical user of the technologies at issue won't regularly encounter misinformation; but the principle doesn't tell us what the specific content of those laws or regulations should be. It's clear that the sorts of policies that some social media companies already employ—such as labelling certain content as “disputed,” or pointing users to more accurate information—are not sufficient to prevent a typical user from regularly encountering misinformation. But, on the other hand, outright bans of misinformation may not be required; instead, policies aimed at diminishing the visibility of misinformation might be more effective at preventing average users from being exposed to such content. In any case, since by establishing laws or regulations designed to prevent social media platforms and search engines from regularly dispensing misinformation we are not violating any competing epistemic obligations, the Shared Obligation Principle entails that we all have a shared epistemic obligation to work together to establish some such laws or regulations.

(One might wonder whether our shared obligation to restrict the distribution of misinformation extends beyond search engines and social media platforms—after all, an unfortunately wide variety of websites and media organizations regularly expose their readers and viewers to misinformation. The foregoing argument focuses on search engines and social media platforms because avoiding such sources is sufficiently difficult that most individuals can't be reasonably expected to do so. Conversely, most individuals will find it sufficiently easy to avoid relying on any particular website, newspaper, or the like. However, while the foregoing argument cannot establish that we have a shared epistemic obligation to restrict all such sources from distributing misinformation, it entails that doing so would be epistemically permissible. While the individual citizens of democratic states will find it sufficiently easy to avoid relying on problematic sources while acting individually, it still might be more convenient for us to fulfill our obligation to avoid unnecessary exposure to misinformation by performing some joint action; adopting laws or regulations that prohibit any media company from regularly distributing misinformation is epistemically permissible, then, because it is a particularly convenient

⁵⁹One might point to a different sort of difficulty: perhaps there is no feasible process by which a government can restrict misinformation because there is no feasible process by which a government can determine what claims constitute misinformation (for discussion of a related, though more general, issue, see [Leiter, 2016, §IV]). However, psychology researchers have a simple and effective method for determining what constitutes misinformation: they rely on widespread agreement amongst independent fact-checking organizations. Some social media companies also use this method to label dubious content. There is no obvious reason why governments could not employ a similar strategy.

⁶⁰The view that individuals who are not aware of the extent of their vulnerability to misinformation are not epistemically obligated to avoid misinformation entails that we cannot derive this conclusion by relying on the Shared Obligation Principle. However, a defender of that view might still endorse the conclusion of the present argument if he maintains that a group sometimes possesses a collective obligation even though all the individual group members don't possess that obligation. For a defense of the latter principle, see Björnsson (2014).

way for us to fulfill our epistemic obligations. By way of analogy, suppose that the individual citizens of some state each have a moral obligation to help the victims of some recent natural disaster. If these citizens have convenient methods for helping while acting individually, we can't conclude that they have a shared obligation to provide help via some joint action; but, if it is particularly convenient for these citizens to fulfill their moral obligation by performing some joint action—for instance, by enabling their government to organize and fund relief efforts—surely it is morally permissible for them to perform some such joint action.)

6 | CONCLUSION

Ultimately, then, there is a purely epistemic justification for restricting the public expression of certain falsehoods. Because we are all vulnerable to false information in systematic and unavoidable ways, each of us has an individual epistemic obligation to avoid relying unnecessarily on sources that regularly dispense misinformation—especially when by relying on a given source we would encounter the same misinformation repeatedly, or encounter misinformation consistent with what we already believe. When we rely on popular internet technologies such as social media platforms and search engines, not only are we likely to encounter misinformation, but we are likely to encounter the same misinformation repeatedly, and that misinformation is likely to be consistent with what we already believe; so, each of us has an epistemic obligation to avoid obtaining information via social media and search engines unnecessarily. But fulfilling this obligation while acting alone is so difficult that most of us can't be reasonably expected to do so; conversely, it is relatively easy for each of us to fulfill our obligation to avoid relying on sources that regularly dispense misinformation by contributing to a corresponding joint action. Specifically, those of us who are citizens of democratic states are able, via joint action, to establish laws or regulations designed to prevent social media platforms and search engines from regularly dispensing misinformation; and consequently, we have a shared epistemic obligation to establish such laws or regulations.⁶¹

REFERENCES

- Ahlstrom-Vij, H. K., & Dunn, J. (2018). Introduction: Epistemic consequentialism. In H. K. Ahlstrom-Vij & J. Dunn (Eds.), *Epistemic consequentialism* (pp. 1–22). Oxford UP.
- Ahlstrom-Vij, K. (2013). *Epistemic paternalism: A defence*. Palgrave Macmillan.
- Alston, W. (1985). Concepts of epistemic justification. *The Monist*, 68, 57–89.
- Anspach, N., & Carlson, T. (2020). What to believe? Social media commentary and belief in misinformation. *Political Behavior*, 42, 697–718.
- Avaaz. (2020a). “How Facebook can Flatten the Curve of the Coronavirus Infodemic.” *Avaaz Report*, April 15, 2020. https://avaazimages.avaaz.org/facebook_coronavirus_misinformation.pdf
- Avaaz. (2020b). “Why is YouTube Broadcasting Climate Misinformation to Millions?” *Avaaz Report*, January 15, 2020. https://avaazimages.avaaz.org/youtube_climate_misinformation_report.pdf

⁶¹For comments that led to significant improvements to this paper, thanks to Neil Levy, Neil Mehta, and an anonymous reviewer for *Analytic Philosophy*.

- Bakshy, E., Messing, S., & Adamic, L. A. (2015). Exposure to ideologically diverse news and opinion on Facebook. *Science*, *348*, 1130–1132.
- Berker, S. (2013). Epistemic teleology and the separateness of propositions. *Philosophical Review*, *122*, 337–393.
- Bhattacharya, A. (2016). Google Has No Intention of Removing a Holocaust-Denier from the Top of Its ‘Did the Holocaust Happen?’ Search Result. *Quartz*, December 13, 2016. <https://qz.com/862285/google-holocaust-denier-search-result-did-the-holocaust-happen/>
- Björnsson, G. (2014). Essentially shared obligations. *Midwest Studies in Philosophy*, *38*, 103–120.
- Bovet, A., & Makse, H. (2019). Influence of fake news in twitter during the 2016 US presidential election. *Nature Communications*, *10*, 7.
- Brashier, N., Eliseev, E., & Marsh, E. (2020). An initial accuracy focus prevents illusory truth. *Cognition*, *194*, 104054.
- Brashier, N., & Marsh, E. (2020). Judging truth. *Annual Review of Psychology*, *71*, 499–515.
- Bronstein, M., Pennycook, G., Bear, A., Rand, D., & Cannon, T. (2019). Belief in fake news is associated with delusionality, dogmatism, religious fundamentalism, and reduced analytic thinking. *Journal of Applied Research in Memory and Cognition*, *8*, 108–117.
- Brown, J. (2020). What is epistemic blame? *Noûs*, *54*, 389–407.
- Bump, P. (2016). “Google’s Top News Link for ‘Final Election Results’ Goes to a Fake News Site with False Numbers.” *The Washington Post*, November 14, 2016. <https://www.washingtonpost.com/news/the-fix/wp/2016/11/14/googles-top-news-link-for-final-election-results-goes-to-a-fake-news-site-with-false-numbers/>
- Bush, D., & Zaheer, A. (2019). “Bing’s Top Search Results Contain an Alarming Amount of Disinformation.” *Stanford Internet Observatory*, December 17, 2019. <https://cyber.fsi.stanford.edu/io/news/bing-search-disinformation>
- Chan, M.-p. S., Jones, C., Jamieson, K., & Albarracín, D. (2017). Debunking: A meta-analysis of the psychological efficacy of messages countering misinformation. *Psychological Science*, *28*, 1531–1546.
- De Keersmaecker, J., Dunning, D., Pennycook, G., Rand, D., Sanchez, C., Unkelbach, C., & Roets, A. (2020). Investigating the robustness of the illusory truth effect across individual differences in cognitive ability, need for cognitive closure, and cognitive style. *Personality and Social Psychology Bulletin*, *46*, 204–215.
- DiFonzo, N., Beckstead, J., Stupak, N., & Walders, K. (2016). Validity judgments of rumors heard multiple times: The shape of the truth effect. *Social Influence*, *11*, 22–39.
- Ditto, P., Liu, B., Clark, C., Wojcik, S., Chen, E., Grady, R., Celniker, J., & Zinger, J. (2019). At least bias is bipartisan: A meta-analytic comparison of partisan bias in liberals and conservatives. *Perspectives on Psychological Science*, *14*, 273–291.
- Ditto, P., & Lopez, D. (1992). Motivated skepticism: Use of differential decision criteria for preferred and nonpreferred conclusions. *Journal of Personality and Social Psychology*, *63*, 568–584.
- Dylko, I., Dolgov, I., Hoffman, W., Eckhart, N., Molina, M., & Aaziz, O. (2017). The dark side of technology: An experimental investigation of the influence of customizability technology on online political selective exposure. *Computers in Human Behavior*, *73*, 181–190.
- Ecker, U., & Ang, L. C. (2019). Political attitudes and the processing of misinformation corrections. *Political Psychology*, *40*, 241–260.
- Fallis, D. (2016). Mis- and dis-information. In L. Floridi (Ed.), *The Routledge handbook of philosophy of information* (pp. 332–346). Routledge.
- Fazio, L., Barber, S., Rajaram, S., & Ornstein, P. (2013). Creating illusions of knowledge: Learning errors that contradict prior knowledge. *Journal of Experimental Psychology: General*, *142*, 1–5.
- Fazio, L., & Marsh, E. (2008). Slowing presentation speed increases illusions of knowledge. *Psychonomic Bulletin & Review*, *15*, 180–185.
- Fazio, L., Nadia Brashier, B., Payne, K., & Marsh, E. (2015). Knowledge does not protect against illusory truth. *Journal of Experimental Psychology: General*, *144*, 993–1002.
- Fazio, L., Rand, D., & Pennycook, G. (2019). Repetition increases perceived truth equally for plausible and implausible statements. *Psychonomic Bulletin & Review*, *26*, 1705–1710.
- Feinberg, J. (1968). Collective responsibility. *The Journal of Philosophy*, *65*, 674–688.
- Feldman, R. (2002). Epistemological duties. In P. Moser (Ed.), *The Oxford handbook of epistemology* (pp. 362–384). Oxford University Press.
- Franklin, C. (2013). A theory of the normative force of pleas. *Philosophical Studies*, *163*, 479–502.
- Gaines, B., & Mondak, J. (2009). Typing together? Clustering of ideological types in online social networks. *Journal of Information Technology & Politics*, *6*, 216–231.

- Gilbert, D., Tafarodi, R., & Malone, P. (1993). You can't not believe everything you read. *Journal of Personality and Social Psychology*, *65*, 221–233.
- Gordon, A., Quadflieg, S., Brooks, J., Ecker, U., & Lewandowsky, S. (2019). Keeping track of 'alternative facts': The neural correlates of processing misinformation corrections. *NeuroImage*, *193*, 46–56.
- Green, M., & Donahue, J. (2011). Persistence of belief change in the face of deception: The effect of factual stories revealed to be false. *Media Psychology*, *14*, 312–331.
- Grinberg, N., Joseph, K., Friedland, L., Swire-Thompson, B., & Lazer, D. (2019). Fake news on twitter during the 2016 U.S. presidential election. *Science*, *363*, 374–378.
- Halberstam, Y., & Knight, B. (2016). Homophily, group size, and the diffusion of political information in social networks: Evidence from twitter. *Journal of Public Economics*, *143*, 73–88.
- Hall, R., & Johnson, C. (1998). The epistemic duty to seek more evidence. *American Philosophical Quarterly*, *35*, 129–139.
- Hasson, U., Simmons, J., & Todorov, A. (2005). Believe it or not: On the possibility of suspending belief. *Psychological Science*, *16*, 566–571.
- Henkel, L., & Mattson, M. (2011). Reading is believing: The truth effect and source credibility. *Consciousness and Cognition*, *20*, 1705–1721.
- Himmelboim, I., McCreery, S., & Smith, M. (2013). Birds of a feather tweet together: Integrating network and content analyses to examine cross-ideology exposure on twitter. *Journal of Computer-Mediated Communication*, *18*, 154–174.
- Hindriks, F. (2019). The duty to join forces: When individuals lack control. *The Monist*, *102*, 204–220.
- Jerit, J., & Zhao, Y. (2020). Political misinformation. *Annual Review of Political Science*, *23*, 77–94.
- Johnson, H., & Seifert, C. (1994). Sources of the continued influence effect: When misinformation in memory affects later inferences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 1420–1436.
- Knobloch-Westerwick, S., Johnson, B., & Westerwick, A. (2015). Confirmation bias in online searches: Impacts of selective exposure before an election on political attitude strength and shifts. *Journal of Computer-Mediated Communication*, *20*, 171–187.
- Kornblith, H. (1983). Justified belief and epistemically responsible action. *The Philosophical Review*, *92*, 33–48.
- Lackey, J. (2021). Epistemic duties regarding others. In K. McCain & S. Stapleford (Eds.), *Epistemic duties: New arguments, New Angles*. Routledge.
- Lapowsky, I. (2018). "Google Autocomplete Still Makes Vile Suggestions." *Wired*, February 12, 2018. <https://www.wired.com/story/google-autocomplete-vile-suggestions/>
- Leiter, B. (2016). The case against free speech. *Sydney Law Review*, *38*, 407–439.
- Leon, M. (2002). Responsible believers. *The Monist*, *85*, 421–435.
- Levy, N. (2012). Ecological engineering: Reshaping our environments to achieve our goals. *Philosophy & Technology*, *25*, 589–604.
- Levy, N. (2017). The bad news about fake news. *Social Epistemology Review and Reply Collective*, *6*, 20–36.
- Levy, N. (2019). No-platforming and higher-order evidence, or anti-anti-no-platforming. *Journal of the American Philosophical Association*, *5*, 487–502.
- Levy, N., & Mandelbaum, E. (2014). The powers that bind: Doxastic voluntarism and epistemic obligation. In J. Matheson & R. Vitz (Eds.), *The ethics of belief: Individual and social* (pp. 15–32). Oxford University Press.
- Lewandowsky, S., Ecker, U., Seifert, C., Schwarz, N., & Cook, J. (2012). Misinformation and its correction: Continued influence and successful debiasing. *Psychological Science in the Public Interest*, *13*, 106–131.
- Mandelbaum, E. (2014). Thinking is believing. *Inquiry*, *57*, 55–96.
- Mathiesen, K. (2019). Fake news and the limits of freedom of speech. In C. Fox & J. Saunders (Eds.), *Media ethics, free speech, and the requirements of democracy* (pp. 161–179). Routledge.
- May, L. (1992). *Sharing responsibility*. University of Chicago Press.
- Mercier, H. (2017). How gullible are we? A review of the evidence from psychology and social science. *Review of General Psychology*, *21*, 103–122.
- Mill, J. S. (1859/1977). On liberty. In J. M. Robson (Ed.), *Collected works of John Stuart Mill* (Vol. 18). University of Toronto Press.
- Millar, B. (2021). Shared epistemic responsibility. *Episteme*, *18*, 493–506.
- Miller, B., & Record, I. (2013). Justified belief in a digital age: On the epistemic implications of secret internet technologies. *Episteme*, *10*, 117–134.

- Miller, B., & Record, I. (2017). Responsible epistemic technologies: A social-epistemological analysis of auto-completed web search. *New Media & Society*, *19*, 1945–1963.
- Miller, S., & Makela, P. (2005). The collectivist approach to collective moral responsibility. *Metaphilosophy*, *36*, 634–651.
- Montmarquet, J. (1993). *Epistemic virtue and doxastic responsibility*. Rowman and Littlefield.
- Nadarevic, L., & Abfal, A. (2017). Unveiling the truth: Warnings reduce the repetition-based truth effect. *Psychological Research*, *81*, 814–826.
- Narveson, J. (2002). Collective responsibility. *The Journal of Ethics*, *6*, 179–198.
- Nickerson, R. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, *2*, 175–220.
- Nottelmann, N. (2007). *Blameworthy belief: A study in epistemic deontologism*. Springer.
- O'Connor, C., & Weatherall, J. (2019). *The misinformation age: How false beliefs spread*. Yale University Press.
- Ohlheiser, A. (2017). "Why Google Once Claimed that Obama was Staging a Coup." *The Washington Post*, March 8, 2017. <https://www.washingtonpost.com/news/the-intersect/wp/2017/03/08/why-google-once-claimed-that-obama-was-staging-a-coup/>
- Olsen, K. (2017). A defense of the objective/subjective moral ought distinction. *The Journal of Ethics*, *21*, 351–373.
- Pantazi, M., Kissine, M., & Klein, O. (2018). The power of the truth bias: False information affects memory and judgment even in the absence of distraction. *Social Cognition*, *36*, 167–198.
- Peels, R. (2017). *Responsible belief: A theory in ethics and epistemology*. Oxford University Press.
- Pennycook, G., Cannon, T., & Rand, D. (2018). Prior exposure increases perceived accuracy of fake news. *Journal of Experimental Psychology: General*, *147*, 1865–1880.
- Pennycook, G., & Rand, D. (2019). Lazy, not biased: Susceptibility to partisan fake news is better explained by lack of reasoning than by motivated reasoning. *Cognition*, *188*, 39–50.
- Pennycook, G., & Rand, D. (2020). Who falls for fake news? The roles of bullshit receptivity, overclaiming, familiarity, and analytic thinking. *Journal of Personality*, *88*, 185–200.
- Porter, E., & Wood, T. (2019). *False alarm: The truth about political mistruths in the trump era*. Cambridge UP.
- Rapp, D. (2016). The consequences of reading inaccurate information. *Current Directions in Psychological Science*, *25*, 281–285.
- Rapp, D., Hinze, S., Kohlhepp, K., & Ryskin, R. (2014). Reducing reliance on inaccurate information. *Memory and Cognition*, *42*, 11–26.
- Rettler, L., & Rettler, B. (2021). Epistemic duty and implicit bias. In K. McCain & S. Stapleford (Eds.), *Epistemic duties: New arguments, new angles*. Routledge.
- Richter, T., Schroeder, S., & Wöhrmann, B. (2009). You don't have to believe everything you read: Background knowledge permits fast and efficient validation of information. *Journal of Personality and Social Psychology*, *96*, 538–558.
- Roberts, R., & Jay Wood, W. (2007). *Intellectual virtues: An essay in regulative epistemology*. Oxford UP.
- Shin, J., Jian, L., Driscoll, K., & Bar, F. (2017). Political rumoring on twitter during the 2012 US presidential election: Rumor diffusion and correction. *New Media & Society*, *19*, 1214–1235.
- Sliwa, P. (2019). The power of excuses. *Philosophy & Public Affairs*, *47*, 37–71.
- Sverdlik, S. (1987). Collective responsibility. *Philosophical Studies*, *51*, 61–76.
- Swire, B., Berinsky, A., Lewandowsky, S., & Ecker, U. (2017). Processing political misinformation: Comprehending the trump phenomenon. *Royal Society Open Science*, *4*, 160802.
- Sylvan, K. (2020). An epistemic nonconsequentialism. *Philosophical Review*, *129*, 1–51.
- Taber, C., & Lodge, M. (2006). Motivated skepticism in the evaluation of political beliefs. *American Journal of Political Science*, *50*, 755–769.
- Thorson, E. (2016). Belief echoes: The persistent effects of corrected misinformation. *Political Communication*, *33*, 460–480.
- Unkelbach, C., & Koch, A. (2019). Gullible but functional? Information repetition and the formation of beliefs. In J. Forgas & R. Baumeister (Eds.), *The social psychology of gullibility* (pp. 42–60). Routledge.
- van der Linden, S., Leiserowitz, A., Rosenthal, S., & Maibach, E. (2017). Inoculating the public against misinformation about climate change. *Global Challenges*, *1*, 1600008.
- Vicario, D., Michela, A. B., Zollo, F., Petroni, F., Scala, A., Guido Caldarelli, H., Stanley, E., & Quattrociochi, W. (2016). The spreading of misinformation online. *PNAS*, *113*, 554–559.

- Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, 359, 1146–1151.
- Walter, N., Jonathan Cohen, R., Holbert, L., & Morag, Y. (2020). Fact-checking: A meta-analysis of what works and for whom. *Political Communication*, 37, 350–375.
- Whitcomb, D., Battaly, H., Baehr, J., & Howard-Snyder, D. (2017). Intellectual humility: Owning our limitations. *Philosophy and Phenomenological Research*, 94, 509–539.
- Zagzebski, L. (1996). *Virtues of the mind: An inquiry into the nature of virtue and the ethical foundations of knowledge*. Cambridge UP.