

# Epistemic Obligations of the Laity

Boyd Millar  
Washington University in St. Louis  
Email: millar.boyd@gmail.com

Forthcoming in *Episteme*

## Abstract

Very often when the vast majority of experts agree on some scientific issue, laypeople nonetheless regularly consume articles, videos, lectures, etc., the principal claims of which are inconsistent with the expert consensus. Moreover, it is standardly assumed that it is entirely appropriate, and perhaps even obligatory, for laypeople to consume such anti-consensus material. I maintain that this standard assumption gets things backwards. Each of us is particularly vulnerable to false claims when we are not experts on some topic—such falsehoods have systematic negative impacts on our doxastic attitudes that we can neither prevent nor correct. So, when there is clear expert consensus on a given scientific issue, while it is permissible for experts to consume anti-consensus material, laypeople have an epistemic obligation to avoid such material. This argument has important consequences for philosophical discussions of our epistemic obligations to perform or omit belief-influencing actions. Such discussions typically abstract away from the important differences between experts and laypeople. Accordingly, we should reject this typical practice as problematic, and insist instead that laypeople and experts have fundamentally different epistemic obligations.

**Keywords:** Epistemic obligations; epistemic consequentialism; epistemic virtues; experts

## 1. Introduction

Suppose you don't know much about climate change beyond the fact that most all climate scientists think the planet is getting warmer largely thanks to human activity; and suppose you come across an article in a newspaper or magazine with a headline along the lines of "The Science of Climate Change is not Settled." Should you read it?

We can interpret this question in purely epistemic terms. According to a standard view, in addition to our *moral* and *prudential* obligations, human beings have a distinct set of *epistemic* obligations.<sup>1</sup> Some of these obligations will be obligations to form or maintain only certain sorts of beliefs (more generally, doxastic attitudes). But, 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 (we can understand

---

<sup>1</sup> For defences of this view, see, for example, Alston (1985), Feldman (2002), Nottelmann (2007, §1.3), Peels (2017, 101-102), and Brown (2020).

epistemic success to be either true or accurate belief or knowledge); and as such, some of our epistemic obligations are obligations to perform or omit certain belief-influencing actions.<sup>2</sup> Understood in these terms, the question at issue is: is it epistemically permissible for you to read the article?

Probably the most natural answer to this question is “*of course it is.*” After all, you won’t be compelled to accept any of the author’s claims; and perhaps it’s best to be familiar with the arguments on both sides of such a controversial and important issue. However, if Levy (2006) is right, it’s highly likely that reading the article will leave you worse off epistemically. For instance, because you don’t know very much about the topic, you won’t be able to spot all of the author’s mistakes. You might end up believing at least some of the author’s false claims; and even those falsehoods that you avoid believing are likely to diminish the accuracy of your beliefs on the topic. Arguably, then, we should agree with Levy that you ought not to read the article.<sup>3</sup>

If you have an epistemic obligation not to read the article, the source of this obligation is plausibly your lack of expertise regarding climate science. For instance, a practicing climate scientist doesn’t have the same reasons to avoid reading the article—she’ll be able to spot the author’s mistakes, and so her doxastic attitudes aren’t likely to be influenced by the author’s false claims. Yet, discussions of our epistemic obligations to perform or omit belief-influencing actions typically abstract away from these important differences between believers: we’re told that we all have an epistemic obligation to gather evidence, or to avoid dubious sources of information, or to diversify the sources of information we rely on.<sup>4</sup> Conversely, cases of the sort just described suggest that experts and laypeople possess fundamentally different epistemic obligations. And if so, it is crucially important that discussions of our epistemic obligations bear the distinction between experts and laypeople in mind—after all, we all have beliefs concerning a great many topics, and yet each of us is an expert regarding very little. It may be that, with respect to topics concerning which we are experts, we ought to behave a certain way (perhaps we ought to conduct thorough investigations and reflect carefully on the evidence we collect); and it

---

<sup>2</sup> For defences of this view, see, for example, Kornblith (1983), Alston (1985), Leon (2002), Nottelmann (2007, chap. 12), Peels (2017, chap. 3), and Lackey (2021).

<sup>3</sup> More specifically, Levy’s (2006, 56 & 60) view is that you ought not to read the article so long as its content is inconsistent with what you already believe (see note 5 below).

<sup>4</sup> As I’ve indicated already, Levy (2006) is an important exception to this rule. Hall and Johnson (1998) maintain that we have an epistemic obligation to gather evidence. Nottelmann (2007, 174-175) maintains that we have an epistemic obligation to avoid dubious sources (Levy and Mandelbaum (2014) defend a qualified form of that thesis). Worsnip (2019) maintains that we have an epistemic obligation to diversify our sources.

may be that, with respect to most other topics, we ought to behave very differently (perhaps our principal obligation is to allocate trust appropriately and consult the right people).

In order to establish that experts and laypeople have fundamentally different epistemic obligations, I will focus on scientific issues where a clear consensus exists amongst the relevant experts: that the climate is changing significantly largely due to human activity, that species result from evolution rather than intelligent design, that vaccines are safe and effective, that COVID-19 is more dangerous than the flu, and so on. With respect to such issues, I maintain that laypeople, but not experts, have an epistemic obligation to avoid unnecessary exposure to *anti-consensus material*: articles, videos, lectures, etc., the principal claims of which are inconsistent with the expert consensus.

Claiming that laypeople have an obligation to actively avoid dissenting voices in such cases might seem to be endorsing a dangerous form of dogmatism; however, the argument for this conclusion is grounded in epistemic humility—in a recognition of our cognitive limitations and vulnerabilities. More specifically, the argument is that each of us is particularly vulnerable to false claims when we are not experts on some topic—such falsehoods have systematic negative impacts on our doxastic attitudes that we can neither prevent nor correct. And while we can't avoid falsehoods generally, there is a particular category of falsehood that we can often avoid: falsehoods that are inconsistent with the available evidence. Since, then, we are likely to encounter false claims inconsistent with the available evidence when we consume anti-consensus material; and since the doxastic attitudes of laypeople, but not experts, are likely to be systemically negatively influenced by encountering such falsehoods; laypeople, but not experts, have an epistemic obligation not to consume anti-consensus material unnecessarily.<sup>5</sup> The conclusion is not that you ought to believe whatever the experts say simply because they say it—the conclusion defended here says nothing about what you ought to believe, and it certainly does

---

<sup>5</sup> The present argument differs from Levy's (2006) argument in a few important respects. First, while Levy's argument concerns "controversial questions" generally, the present conclusion is restricted to scientific issues concerning which clear expert consensus exists. Second, Levy's argument is that gathering evidence threatens the justification of laypeople's beliefs and might require one to suspend belief concerning all controversial questions; whereas the present argument will focus largely on the accuracy of one's doxastic attitudes. And third, Levy maintains that you are obligated to avoid material that is inconsistent with what you already believe (so, e.g., a climate science sceptic ought to avoid reading IPCC reports); conversely, I maintain that you are obligated to avoid anti-consensus material regardless of what you already believe. I should also note that Fantl (2018, chaps. 6-7) defends a related conclusion: that laypeople should sometimes avoid engaging with counterarguments to propositions they know. However, Fantl doesn't focus on epistemic obligations as such, and there is very little overlap between his arguments and those presented here.

not imply that you shouldn't study the evidence and arguments supporting the expert consensus. Rather, the point is simply that when you encounter someone who denies the clear expert consensus on some scientific issue, you shouldn't listen to that person.

## **2. Clear expert consensus**

Before proceeding to the argument, I should clarify some points regarding the notion of clear expert consensus. According to most philosophers' definitions, experts are different from laypeople in two chief respects: knowledge and ability.<sup>6</sup> That is, an expert on some topic is someone who possesses significantly greater knowledge of the available evidence relevant to that topic than most people do, and who is significantly better able to interpret and draw conclusions from that evidence than most people are. So understood, expertise comes in degrees—for instance, you might know significantly more than most people on some topic because you took a single University course. But for present purposes, the experts who matter are those who occupy the far end of the spectrum. Consequently, I will assume that someone is an expert on some topic if and only if, first, he knows about as much as anyone concerning the available evidence relevant to that topic, and second, he is about as good as anyone at interpreting and drawing conclusions from that evidence. For instance, with respect to specific scientific topics, the experts are individuals with advanced degrees in some relevant scientific discipline, and who publish research on the topic in peer-reviewed venues.<sup>7</sup>

Since, for a given scientific issue, the relevant experts will be scientists who publish research on related scientific topics, there is clear expert consensus on a given scientific issue just in case there is a large group of relevant scientists who overwhelmingly agree on that issue—that is, the vast majority endorse some particular verdict with respect to that issue. We can understand agreeing about or endorsing some verdict in terms of what these scientists have published, what they believe, or both. So, for instance, given that roughly 85% of the relevant scientists who express a view on the issue when surveyed maintain that anthropogenic climate change is occurring, and given that more than 95% of peer-reviewed papers that take a stand on

---

<sup>6</sup> See, for example, Goldman (2001, 91-92), Anderson (2011, 145), and Ballantyne (2019, chap. 8).

<sup>7</sup> In other words, the experts are individuals who occupy categories f-h in Anderson's (2011, 146-147) "hierarchy of expertise."

the question support that thesis, there is clear expert consensus that the climate is changing largely due to human activity.<sup>8</sup>

In addition to these terminological matters, there are two substantive issues concerning expert consensus that I should address. First, the ensuing argument will assume that those who reject the clear expert consensus on any given scientific issue are more likely to be mistaken than those who endorse it; but one might worry about the reliability of expert consensus. For instance, drawing on Goldman (2001), one might argue that widespread agreement amongst experts does not improve their chances of being correct when those experts have not formed their beliefs independently of one another. However, scientists typically form their beliefs at least partially independently of one another.<sup>9</sup> And, as Coady (2006) and Lackey (2013) show, even if they didn't, the fact that the vast majority of scientists endorse some scientific proposition would still make it more likely that that proposition is true. Accordingly, we ought to insist that someone who endorses the clear expert consensus on some scientific question is significantly more likely to be correct than someone who rejects it: an individual scientist's verdict on some scientific issue is more likely to be correct than an individual non-scientist's verdict on that issue (since the scientist's verdict is more likely to be supported by the available evidence, and claims that are supported by the available evidence are more likely to be true than claims that are not); and a large group of scientists' verdict on some scientific issue is significantly more likely to be correct than a much smaller group of scientists' verdict on that issue (since a large group of scientists' verdict is more likely to be supported by the available evidence than is a much smaller group of scientists' verdict).

Second, the ensuing argument will assume that, at least very often, laypeople are able to determine when clear expert consensus exists on some scientific issue; but one might worry that laypeople face significant difficulties determining what the vast majority of experts believe about any given scientific issue. First, it might be difficult for many laypeople to recognize who the experts are—specifically, it might be difficult to recognize that scientists are the experts on any given scientific question. For instance, evidence of bias, problematic incentives, intellectual dishonesty, or even outright corruption within the scientific community might lead many laypeople to conclude that scientists are not the most reliable interpreters of the evidence

---

<sup>8</sup> See Cook et al. (2013) and Verheggen et al. (2014).

<sup>9</sup> As Goldman (2001, 103) acknowledges.

concerning issues such as climate change and evolution—but that some group of non-scientists, such as political or religious leaders, are better able to draw conclusions from that evidence. However, while there is good evidence that scientists are sometimes biased or dishonest when they investigate scientific questions (especially politically-loaded scientific questions), there is vastly more evidence of bias and dishonesty amongst political and religious leaders (and any other relevant group of non-scientists). Moreover, while there is extensive evidence that scientists are reliable interpreters of the available evidence on a wide range of scientific questions—most all of us encounter examples of successful scientific research on a more or less constant basis—there is very little evidence that any group of non-scientists is similarly reliable. So, it’s not the case that misleading evidence makes it difficult for laypeople to recognize that scientists are the individuals best able to answer even politically-loaded scientific questions.

(Following Levy (2019; forthcoming), one might object that many laypeople have reasonable grounds not to trust the scientific consensus on many scientific issues—namely, that most scientists do not belong to the religious and political groups with which these laypeople identify, and so are not likely to share their values.<sup>10</sup> The suggestion is that if scientists don’t share my religious and political affiliations then they are less likely to be benevolent towards me, and so it’s reasonable for me not to trust their judgement on politically-loaded scientific issues. However, while it might be reasonable to assume that individuals who share my religious and political affiliations are more likely to be benevolent towards me, and so are less likely to deceive or exploit me, it is not reasonable to assume that such individuals are better able than most to answer scientific questions accurately.<sup>11</sup> So, if I were to conclude that, say, the political and religious leaders whom I trust are better able to interpret and draw conclusions from the available evidence concerning climate change and evolution, I would be weighing evidence of benevolence significantly more heavily than evidence of competence—and that’s unreasonable.)

Alternatively, even if it isn’t difficult for laypeople to determine that the scientists are experts, they might find it difficult to determine what the vast majority of experts believe concerning any given scientific issue. For instance, a layperson who regularly encounters scientists asserting that vaccines are safe and effective, might also sometimes encounter self-professed experts asserting that vaccines cause autism—self-professed experts who, while not

---

<sup>10</sup> Rini (2017) defends a related thesis.

<sup>11</sup> Worsnip (2019, §3) makes some similar points when discussing Rini’s view.

scientists, have credentials sufficient to make it seem as though they have a scientist's knowledge and abilities. As Ballantyne (2019, chap. 9) and Levy (forthcoming, chap. 4) emphasize, it can be quite difficult for a layperson to determine whether such a self-professed expert is a genuine expert; however, at least in very many cases, encountering self-professed experts won't make it difficult for a layperson to determine what the vast majority of experts believe. If I determine that the scientists researching relevant topics are all experts on a given scientific issue (which, again, is not particularly difficult), then even if I mistakenly assume that the self-professed experts I encounter are also experts, these individuals will still constitute a tiny subset of all the experts I recognize.<sup>12</sup> As such, there is a simple method I can employ to determine what the vast majority of experts believe concerning some scientific issue: determine what the vast majority of scientists believe concerning that issue. And fortunately, in very many cases, it is not particularly difficult for a layperson to determine what the vast majority of scientists believe concerning some scientific issue. For instance, Anderson (2011, 149) outlines some simple methods: you can read review articles in scientific publications; you can find surveys of scientists working on the topic; or you can read reports or consensus statements (or summaries thereof) from scientific organizations, such as the National Academy of Sciences. Even more simply, you can read books and articles written by scientists and science journalists for a general audience (or watch documentaries, listen to podcasts, etc.): such material will make clear what issues the relevant scientific community considers settled, and what issues remain controversial. Of course, there will be cases where laypeople will find it difficult to identify clear expert consensus when it exists; the present point is just that there are very many important cases where it is not difficult for laypeople to identify clear expert consensus.

### **3. Our vulnerability to falsehoods**

The suggestion that we have an epistemic obligation to avoid reading or listening to certain material will strike many as implausible. This reaction is grounded in the natural assumption that each of us has considerable control over how we respond to what we read and hear: we are

---

<sup>12</sup> I should emphasize most laypeople who reject the expert consensus on issues like climate change, evolution, and vaccine safety, don't do so because they've encountered so many pseudo-experts that they're unable to determine what most experts believe. Rather, most such individuals deny climate change because they trust their political leaders *rather than* scientists; they reject evolution because they trust their religious leaders *rather than* scientists; and they deny that vaccines are safe because they trust activists and parents with terrifying stories *rather than* scientists. (For some evidence supporting these claims, see §3, and note 20.)

free to reject claims we find implausible, to suspend belief whenever we're unsure about some claim, and when we happen to be taken in by some falsehood, we can always correct our mistake with a bit of research. However, considerable psychological research has established that this natural assumption is mistaken. This research shows that, first, the falsehoods we encounter typically have a systematic negative influence on our doxastic attitudes, particularly when we don't have much relevant background knowledge; and second, that there is nothing we can do to either eliminate a falsehood's negative impact when we first encounter it, or to largely correct its negative impact after the fact.<sup>13</sup>

When we consume written or verbal material, the principal mechanism by which we can prevent false claims from influencing our doxastic attitudes is by comparing each claim to our store of existing beliefs. Researchers sometimes refer to this process as *plausibility checking*: each of us checks the claims we encounter against what we already believe in a spontaneous, automatic fashion; and we typically reject claims inconsistent with what we already believe automatically.<sup>14</sup> But, of course, we will only manage to reject a given falsehood using this procedure so long as it is inconsistent with our existing beliefs.

Crucially, when we encounter claims the truth value of which we don't know (for instance, claims concerning which our background beliefs entail no verdict), each of us exhibits a *truth bias*—a tendency to accept such claims as true.<sup>15</sup> This tendency is sufficiently strong that we often accept claims of unknown truth value even when we have very good reasons not to do so. A well-known experiment—Gilbert et al. (1993)—provides particularly striking evidence of this phenomenon. Subjects were presented with a series of statements regarding a fictional crime, and then asked to rate the perpetrator's dangerousness and to recommend a prison sentence. These statements were explicitly labelled as either true or dubious—subjects were

---

<sup>13</sup> Given the well-publicized “replication crisis” associated with social psychology, I should say something to address natural worries concerning the strength of the evidence that will be presented in the present section. First, as the review articles and meta-analyses cited below demonstrate, the psychological phenomena to which the present argument appeals are supported by a wide variety of experimental evidence. I have tried to avoid relying on individual experiments as much as possible (when I describe individual experiments I do so primarily for the sake of illustration). Second, I have also tried to rely on recently published results as much as possible (research methods have improved significantly in recent years—for instance, employing preregistration and relying on larger sample sizes—largely in response to worries arising from the “replication crisis”).

<sup>14</sup> See Richter et al. (2009). Hasson, Simmons, and Todorov (2005) also provide evidence that we can automatically reject claims that are both informative when false and explicitly labelled as false. For a review of the evidence concerning plausibility checking, see Mercier (2017, 104-105)

<sup>15</sup> For a review of the relevant evidence, see Brashier and Marsh (2020, 501-502).



informed that the true statements were taken from the actual crime report while the dubious statements were taken from entirely unrelated crime reports.<sup>16</sup> Gilbert and colleagues found that, when required to perform a mildly demanding task while reading the statements, subjects tended to judge the perpetrator more severely when the dubious statements they had read increased the severity of the crime. However, a more recent version of this experiment—Pantazi et al. (2018)—found that subjects demonstrated a significant tendency to treat the dubious claims as true *even when not subjected to distractions of any kind*.<sup>17</sup>

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*.<sup>18</sup> So, because each separate time we encounter some claim increases our confidence that that claim is true, simply repeating some falsehood often enough can sometimes cause us to believe it.<sup>19</sup> Again, our background beliefs can prevent us from accepting some falsehood, even if it is repeated frequently; but when the claim isn't in tension with our existing beliefs, the effect is remarkably robust. For instance, Henkel and Mattson (2011) showed that individuals who read a particular statement on three separate occasions over the course of two weeks, judged that statement to be either probably or certainly true almost 70% of the time—despite the fact that they were given explicit advance warning that the source of the information was unreliable.

We are also particularly likely to believe false claims made by individuals we trust. As Levy (2019; forthcoming) argues, because human beings are social creatures with limited resources, we have a significant tendency to defer to individuals who belong to groups we identify with, and to individuals who have acquired a certain level of prestige. For instance, a recent study by Barber and Pope (2019) found that Republicans were much more likely to endorse a stereotypically liberal policy when told that Donald Trump supports it. Such deference may sometimes be rational (Levy maintains it typically is); but a consequence of this strategy is that when individuals who belong to the relevant groups (especially prestigious individuals) make false claims that are not in tension with our existing beliefs, there's a good chance we'll

---

<sup>16</sup> Ostensibly, the latter statements were labelled *false*. However, because subjects were told that the “false” statements were taken from unrelated crime reports, they didn't actually know whether these statements were true or false. See Gilbert et al. (1993, 223).

<sup>17</sup> Mercier (2020, chap. 3) sometimes seems to suggest that human beings have a tendency to disbelieve or doubt statements of unknown truth value; however, none of the research he cites supports that specific conclusion, and he doesn't consider the research reviewed in the present paragraph.

<sup>18</sup> For a review of the relevant evidence, see Unkelbach and Koch (2019).

<sup>19</sup> DiFonzo et al. (2016).

end up believing them. For instance, there is considerable evidence that Americans' views about the existence (or seriousness) of climate change has been driven largely by the claims communicated by partisan elites; as a result, large numbers of Americans have acquired false beliefs on the topic.<sup>20</sup>

It's extremely important to recognize that false claims have significant negative impacts on our doxastic attitudes even when we don't believe them. Perhaps most problematically, encountering falsehoods frequently undercuts the positive impact that encountering truths would otherwise have. Being exposed to significant quantities of dubious information can have a paralyzing effect, leading us to conclude that there's no way to tell what's true—a fact that politicians and industry groups often take advantage of. But even a single encounter with false information can have significant effects. For instance, van der Linden et al. (2017) showed that, typically, reading the statement that “97% of climate scientists have concluded that human-caused climate change is happening” significantly improved the accuracy of an individual's estimate of the current level of scientific consensus regarding climate change. However, when such an individual subsequently read the false statement that “there is no consensus on human-caused climate change,” encountering this falsehood *completely eliminated* the accuracy-improving impact that reading the true statement would otherwise have had.

One might think that there are strategies we can employ to protect our doxastic attitudes from the negative influence of false claims; however, when we don't have much in the way of background beliefs concerning some topic, we won't be able to eliminate (only partially mitigate) the negative influence of the falsehoods we encounter. For instance, you might think it would help to read slowly and carefully; or, perhaps it would be better to adopt a sceptical mindset, or to actively develop certain cognitive skills. But research has shown that reading slowly and carefully does not reduce the consequences of encountering false claims.<sup>21</sup> And research has also shown that individual differences in cognitive ability and personality have little or no bearing on one's susceptibility to some of the cognitive features that help make us vulnerable to falsehoods.<sup>22</sup> Alternatively, you might think that it would help to recognize and try

---

<sup>20</sup> See, for example, Carmichael and Brulle (2017), Tesler (2018), and Merkley and Stecula (2018).

<sup>21</sup> See Rapp (2016).

<sup>22</sup> See, for example, DiFonzo et al. (2016), Mercier (2017), and De keersmaecker et al. (2020).

to actively resist the truth bias and the truth effect; but the available evidence suggests that such attempts would produce only minor improvements.<sup>23</sup>

Even so, one might think we can always correct some falsehood's negative influence after the fact—for instance, by fact-checking the dubious claims we encounter. However, research has shown that the negative influence of false claims cannot be largely reversed or corrected. In some cases where you encounter a falsehood and later encounter a correction, you might simply reject the correction out of hand—an outcome that is particularly likely when the correction is incompatible with your deeply held beliefs.<sup>24</sup> More commonly, encountering the correction will improve the accuracy of your doxastic attitudes—but only partially and temporarily. Research has shown that individuals who read a correction of some falsehood 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 falsehood in the first place.<sup>25</sup> Moreover, research has also shown that this beneficial effect of reading a correction 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.”

In a best case scenario you might discover a correction to some false information and find the correction fully convincing. However, even in such cases the false information that you previously encountered but now reject will continue to influence your doxastic attitudes in systematic ways—a phenomenon known as the *continued influence effect*.<sup>26</sup> For instance, Green and Donahue (2011) found that reading a story had the very same influence on the “story-relevant beliefs” of subjects who were later told it was intentionally fabricated as it did on the beliefs of subjects who hadn't been told that it was fabricated.<sup>27</sup> Gordon et al. (2019) recently attempted to determine the mechanism behind this phenomenon via neuroimaging. They found that when you encounter a correction of false information, rather than simply being deleted from the brain, the false information remains stored in memory and continues to be activated when you consider related topics.

---

<sup>23</sup> See Nadarevic and Aßfalg (2017). For discussion of why personal debiasing strategies are very unlikely to succeed, see Levy (2012, 597-598) and Ahlstrom-Vij (2013, §1.4).

<sup>24</sup> See Ecker and Ang (2019).

<sup>25</sup> See Walter et al. (2020).

<sup>26</sup> For a review relevant evidence, see Lewandowsky et al. (2012).

<sup>27</sup> See also Thorson (2016).

#### 4. Laypeople's epistemic obligation

So, when we encounter false claims concerning topics we don't know much about, they are likely to negatively influence our doxastic attitudes in ways we can't prevent or correct. Consequently, it seems plausible that, when possible, we ought to take steps to avoid being exposed to such falsehoods—for instance, plausibly you ought not to seek the testimony of a pathological liar if you can help it. Of course, we can't avoid being exposed to falsehoods generally; but there are certain categories of falsehoods that we can avoid. In particular, we can often avoid falsehoods that are inconsistent with the available evidence. As we've seen, when clear expert consensus exists on some scientific issue, anti-consensus material is significantly more likely to contain falsehoods than is pro-consensus material. As such, we can minimize the false claims concerning such issues that we encounter by avoiding anti-consensus material.

Accordingly, the conclusion defended in the present section is that, when clear expert consensus exists on some scientific issue, laypeople, but not experts, have an epistemic obligation to avoid unnecessary exposure to anti-consensus material.<sup>28</sup> 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 laypeople, but not experts, have an epistemic obligation to avoid unnecessary exposure to anti-consensus material.

First, one might characterize what makes a belief-influencing action epistemically obligatory in consequentialist terms. According to this theory, some action is an epistemic obligation if and only if it leads to epistemically good outcomes—for example, forming true beliefs and avoiding false beliefs, or acquiring knowledge and avoiding ignorance. If you are a

---

<sup>28</sup> Since expertise comes in degrees, I will use “laypeople” to refer to people who are neither experts nor near-experts; most all of us are laypeople regarding most scientific issues in this sense. I also intend the conclusion to be restricted to issues where it's possible for laypeople to identify the clear expert consensus; however, I will leave this qualification implicit going forward.

layperson concerning some scientific issue and you consume anti-consensus material, the falsehoods you encounter are likely to produce a wide range of epistemically bad outcomes. Since you don't have many background beliefs about the topic, you are likely to end up believing the falsehoods you encounter thanks to the truth bias and truth effect; and you will be particularly likely to accept these falsehoods if the author or speaker is a prestigious member of a group with which you identify. Crucially, those falsehoods that you manage to avoid believing will still counteract the beneficial impact that encountering true claims would otherwise have. Moreover, these negative consequences are not likely to be reversed: if you later encounter a correction of some falsehood (and don't automatically 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 falsehood will remain stored in memory and infect the beliefs you form on the topic. Of course, any anti-consensus material you consume might well have many true things to say; but you will be able to learn most all of those truths by consuming pro-consensus material on the topic. And while some anti-consensus material will include truths you can't find elsewhere, the positive epistemic impact of consuming such material will be significantly outweighed by all of the negative impacts we've just reviewed.

Conversely, when an expert consumes anti-consensus material, she is likely to avoid most all of the epistemically bad outcomes that you would likely suffer. Because the expert has extensive background knowledge on the topic, she will automatically reject most of the falsehoods she encounters via plausibility checking. And because she will be better than laypeople at determining who is competent to speak on the topic, she won't have a tendency to accept the claims of prestigious non-experts. Moreover, the expert can employ strategies to protect herself that are not available to laypeople—for instance, research has shown that individuals can largely avoid being influenced by false statements when they actively correct or explicitly rate the accuracy of those statements as they encounter them (a strategy that requires that you know in advance that the statements are false).<sup>29</sup> And whereas the expert isn't likely to suffer the epistemic harms that a layperson would when consuming anti-consensus material, she is likely to achieve the epistemic benefits—particularly when some anti-consensus material is produced by a genuine expert, there is a good chance that the expert who consumes it will learn something important that she wouldn't have learned otherwise. So, for an expert, consuming

---

<sup>29</sup> See Rapp et al. (2014) and Brashier, Eliseev, and Marsh (2020).

anti-consensus material is not likely to do much harm and will frequently produce significant epistemic benefits.

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. For instance, 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. (Another way to express the point—drawing on Sylvan (2020)—would be that our actions should always manifest respect for truth or accuracy.) The characteristic feature of a Kantian view is that it prohibits the sorts of epistemic trade-offs that a consequentialist view permits—for instance, the Kantian insists that you ought not to acquire false beliefs in order to acquire other true beliefs. So, in order to act in accordance with this Kantian theory you must do everything you can to acquire only true beliefs; and in cases where clear expert consensus exists on some scientific issue, that means, for a layperson, consuming only pro-consensus material. If you, a layperson, consume anti-consensus material because you think you might learn something important, you're making a trade-off: you're allowing yourself to form a number of false (or less accurate) beliefs because there is a chance that you might acquire some true beliefs. (Another way to express the point: given that a layperson's doxastic attitudes are likely to be negatively influenced in systematic ways when she, say, reads an article dismissing climate change, or listens to a speech denying evolution, her actions do not manifest respect for truth or accuracy when she does these things.) Conversely, because an expert isn't likely to acquire false (or less accurate) beliefs when he consumes anti-consensus material, he is not making a similar trade-off when he does so; accordingly, the expert doesn't violate the epistemic analogue of Kant's categorical imperative when he consumes anti-consensus material.

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 only if it is an action that an epistemically virtuous person would characteristically perform. A standard assumption is that an epistemically virtuous person is one who is motivated by the desire to acquire true beliefs (or knowledge) and avoid false beliefs.<sup>30</sup> Such an individual will not perform actions likely to diminish the accuracy of her doxastic

---

<sup>30</sup> See, for example, Montmarquet (1993) and Zagzebski (1996).

attitudes unnecessarily. For instance, an epistemically virtuous person is epistemically cautious and possesses epistemic humility.<sup>31</sup> So, she is aware that she is particularly vulnerable to falsehoods when she has little background knowledge on some topic; and consequently, she takes the only precaution that will ensure that she maintains true beliefs and avoid false beliefs—namely, she avoids anti-consensus material. For a layperson to consume anti-consensus material because she believes she isn't likely to be influenced by the falsehoods it contains is to exhibit epistemic overconfidence; and to consume such material because she finds it reassuring, or engaging, or because she simply doesn't care about the risks, is to exhibit epistemic carelessness.

One might think that an open-minded individual will seek out anti-consensus material—that to always avoid such material is to exhibit the epistemic vice of dogmatism. But since epistemic virtues are directed at truth, being open-minded only requires giving some author or speaker a fair hearing when doing so has a reasonable chance of helping one form true beliefs and avoid false beliefs.<sup>32</sup> Accordingly, open-mindedness would only require that laypersons consume anti-consensus material so long as they could do so without being negatively influenced by the falsehoods such material is likely to contain. Cassam (2019, chap. 5) suggests that a self-confident thinker will trust himself not to be taken in by misleading evidence; but, again, a layperson who isn't worried about being influenced by the falsehoods contained in anti-consensus material is inappropriately overconfident. A layperson who avoids anti-consensus material because he recognizes that his lack of background knowledge and various cognitive limitations make him vulnerable to the falsehoods such material contains, exhibits epistemic humility, not dogmatism.

Conversely, an epistemically virtuous expert will sometimes consume anti-consensus material. Because an expert isn't vulnerable to the falsehoods anti-consensus material contains, his desire to acquire true beliefs and avoid false beliefs won't motivate him to avoid such material. Neither does he exhibit epistemic overconfidence or carelessness if he assumes that he isn't likely to be influenced by the falsehoods such material contains—that assumption is entirely appropriate given his store of relevant background knowledge. In fact, it seems clear that an open-minded expert will sometimes consume anti-consensus material—in particular, when that material is produced by a genuine expert. An expert who refuses to consider expert-generated

---

<sup>31</sup> See, for example, Whitcomb et al. (2017).

<sup>32</sup> See Kwong (2017, 1620).

anti-consensus material, even though there is little risk he will be taken in by any false claims it contains, thereby rejects the opportunity to possibly learn something novel and important. So, an expert who maintains that no anti-consensus material is ever worth considering exhibits dogmatism.

(A potential objection to the preceding argument is that most laypeople aren't obligated to avoid unnecessary exposure to anti-consensus material because most laypeople 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. Alternatively, we could respond to this objection by appealing to the distinction between *objective* and *subjective* obligations<sup>33</sup>—that is, we could maintain that laypeople who are aware of their vulnerability to false information have both a subjective and an objective obligation to avoid unnecessary exposure to anti-consensus material, while laypeople who are ignorant of this vulnerability have only an objective obligation to avoid unnecessary exposure to anti-consensus material. However, some philosophers will insist that there is no sense in which a layperson has an epistemic obligation to avoid anti-consensus material if she has no reason to believe (from an internalist perspective) that she is vulnerable to false information in the ways at issue. Such philosophers should interpret the present argument as establishing a conditional thesis: *if* a layperson is aware of her vulnerability to false claims, *then* she has an epistemic obligation to avoid unnecessary exposure to anti-consensus material.<sup>34</sup>)

## 5. Conclusion

Consequently, each of the most plausible theories of epistemically obligatory action entails that laypeople, but not experts, have an epistemic obligation to avoid unnecessary exposure to anti-consensus material. The overarching principle here is that each of us possesses cognitive features that make us vulnerable to falsehoods—particularly falsehoods that concern topics regarding which we know very little. As such, we ought to try to avoid being exposed to such falsehoods when we can. And when we are laypeople regarding some scientific issue where

---

<sup>33</sup> For discussion, see, for example, Olsen (2017).

<sup>34</sup> Levy and Mandelbaum (2014) defend a related thesis.



clear expert consensus exists, we have an available method for minimizing the falsehoods concerning that topic that we encounter: we can refuse to consume any anti-consensus material (since it is significantly more likely than pro-consensus material to contain falsehoods). But, of course, if we happen to be an expert on some scientific issue, we don't run the same risks when we consume anti-consensus material; so, while experts may, laypeople ought not to consume anti-consensus material unnecessarily. We've now seen, then, that experts and laypeople have fundamentally different epistemic obligations; accordingly, these important differences between believers ought to shape philosophical examinations of our epistemic obligations to perform or omit belief-influencing actions.

The widespread consumption of anti-consensus material has caused enormous numbers of laypeople to form incredibly dangerous false beliefs concerning topics such as climate change, evolution, vaccines, and COVID-19. Fortunately, then, the foregoing argument enables us to give everyone some rather specific advice regarding how best to conduct one's inquiries.<sup>35</sup> If you are a layperson regarding a scientific issue concerning which there is clear expert consensus, then if you encounter an author or speaker whose principal claims are advertised to be inconsistent with that consensus, don't read that author or listen to that speaker. For instance, if a television pundit starts going on about how COVID-19 is no more dangerous than the flu, or how it can be easily treated with hydroxychloroquine or colloidal silver, change the channel. If a friend or relative has some unconventional views about the dangers of vaccines, don't let him share them with you. If a leader of your religious community is going to be speaking about the lie that is the theory of evolution, don't attend services that day. And if you come across an article or book claiming that anthropogenic climate change isn't a serious problem—regardless of whether the author is a prestigious political figure, a respected intellectual, or even a legitimate climate scientist—make sure you don't read it.<sup>36</sup>

---

<sup>35</sup> Giving such advice is sometimes referred to as *regulative epistemology*. For an overview of the topic, see Ballantyne (2019, chaps. 1-2).

<sup>36</sup> For comments that lead to significant improvements to this paper, my thanks to Neil Levy and an anonymous reviewer for *Episteme*.

## References

- Ahlstrom-Vij, Kristoffer. 2013. *Epistemic Paternalism: A Defence*. New York: Palgrave Macmillan.
- Alston, William. 1985. "Concepts of Epistemic Justification." *The Monist* 68: 57-89.
- Anderson, Elizabeth. 2011. "Democracy, Public Policy, and Lay Assessments of Scientific Testimony." *Episteme* 8: 144-164.
- Ballantyne, Nathan. 2019. *Knowing Our Limits*. New York: Oxford University Press.
- Barber, Michael, and Jeremy Pope. 2019. "Does Party Trump Ideology? Disentangling Party and Ideology in America." *American Political Science Review* 113: 38-54.
- Brashier, Nadia, Emmaline Eliseev, and Elizabeth Marsh. 2020. "An Initial Accuracy Focus Prevents Illusory Truth." *Cognition* 194: 104054.
- Brashier, Nadia, and Elizabeth Marsh. 2020. "Judging Truth." *Annual Review of Psychology* 71: 499-515.
- Brown, Jessica. 2020. "What is Epistemic Blame?" *Noûs* 54: 389-407.
- Carmichael, Jason, and Robert Brulle. 2017. "Elite Cues, Media Coverage, and Public Concern: an Integrated Path Analysis of Public Opinion on Climate Change, 2001–2013." *Environmental Politics* 26: 232-252.
- Cassam, Quassim. 2019. *Vices of the Mind: From the Intellectual to the Political*. Oxford: Oxford University Press.
- Coady, David. 2006. "When Experts Disagree." *Episteme* 3: 68-79.
- Cook, John, Dana Nuccitelli, Sarah Green, Mark Richardson, Bärbel Winkler, Rob Painting, Robert Way, et al. 2013. "Quantifying the Consensus on Anthropogenic Global Warming in The Scientific Literature." *Environmental Research Letters* 8: 024024.
- De keersmaecker, Jonas, David Dunning, Gordon Pennycook, David Rand, Carmen Sanchez, Christian Unkelbach, and Arne Roets. 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, Nicholas, Jason Beckstead, Noah Stupak, and Kate Walders. 2016. "Validity Judgments of Rumors Heard Multiple Times: The Shape of the Truth Effect." *Social Influence* 11: 22-39.
- Ecker, Ullrich, and Li Chang Ang. 2019. "Political Attitudes and the Processing of

- Misinformation Corrections.” *Political Psychology* 40: 241-260.
- Fantl, Jeremy. 2018. *The Limitations of the Open Mind*. Oxford: Oxford University Press.
- Feldman, Richard. 2002. “Epistemological Duties.” In *The Oxford Handbook of Epistemology*, edited by Paul Moser, 362-384. New York: Oxford University Press.
- Gilbert, Daniel, Romin Tafari, and Patrick Malone. 1993. “You Can’t Not Believe Everything You Read.” *Journal of Personality and Social Psychology* 65: 221-233.
- Goldman, Alvin. 2001. “Experts: Which Ones Should You Trust?” *Philosophy and Phenomenological Research* 63: 85-110.
- Gordon, Andrew, Susanne Quadflieg, Jonathan Brooks, Ullrich Ecker, and Stephan Lewandowsky. 2019. “Keeping Track of ‘Alternative Facts’: The Neural Correlates of Processing Misinformation Corrections.” *NeuroImage* 193: 46-56.
- Green, Melanie, and John Donahue. 2011. “Persistence of Belief Change in the Face of Deception: The Effect of Factual Stories Revealed to Be False.” *Media Psychology* 14: 312-331.
- Hall, Richard, and Charles Johnson. 1998. “The Epistemic Duty to Seek More Evidence.” *American Philosophical Quarterly* 35: 129-139.
- Hasson, Uri, Joseph Simmons, and Alexander Todorov. 2005. “Believe It or Not: On the Possibility of Suspending Belief.” *Psychological Science* 16: 566-571.
- Henkel, Linda, and Mark Mattson. 2011. “Reading is Believing: The Truth Effect and Source Credibility.” *Consciousness and Cognition* 20: 1705-1721.
- Kornblith, Hilary. 1983. “Justified Belief and Epistemically Responsible Action.” *The Philosophical Review* 92: 33-48.
- Kwong, Jack. 2017. “Is Open-Mindedness Conducive to Truth?” *Synthese* 194: 1613-1626.
- Lackey, Jennifer. 2013. “Disagreement and Belief Dependence: Why Numbers Matter.” In *The Epistemology of Disagreement: New Essays*, edited by David Christensen and Jennifer Lackey, 243-268. Oxford: Oxford University Press.
- . 2021. “Epistemic Duties Regarding Others.” In *Epistemic Duties: New Arguments, New Angles*, edited by Kevin McCain and Scott Stapleford, 281-295. New York: Routledge.
- Leon, Mark. 2002. “Responsible Believers.” *The Monist* 85: 421-435.
- Levy Neil. 2006. “Open-Mindedness and the Duty to Gather Evidence.” *Public Affairs Quarterly* 20: 55-66.

- . 2012. “Ecological Engineering: Reshaping Our Environments to Achieve Our Goals.” *Philosophy & Technology* 25: 589-604.
- . 2019. “Due Deference to Denialism: Explaining Ordinary People’s Rejection of Established Scientific Findings.” *Synthese* 196: 313-327.
- . Forthcoming. *Bad Beliefs: Why They Happen to Good People*. Oxford UP.
- Levy, Neil and Eric Mandelbaum. 2014. “The Powers that Bind: Doxastic Voluntarism and Epistemic Obligation.” In *The Ethics of Belief: Individual and Social*, edited by Jonathan Matheson and Rico Vitz, 15-32. Oxford: Oxford University Press.
- Lewandowsky, Stephan, Ullrich Ecker, Colleen Seifert, Norbert Schwarz, and John Cook. 2012. “Misinformation and Its Correction: Continued Influence and Successful Debiasing.” *Psychological Science in the Public Interest* 13: 106-131.
- Mercier, Hugo. 2016. “Confirmation Bias—Myside Bias.” In *Cognitive Illusions: Intriguing Phenomena in Thinking, Judgment, and Memory*, 2nd edition, edited by Rüdiger Pohl, 99-114. London: Routledge.
- . 2017. “How Gullible Are We? A Review of the Evidence from Psychology and Social Science.” *Review of General Psychology* 21: 103-122.
- . 2020. *Not Born Yesterday: The Science of Who We Trust and What We Believe*. Princeton: Princeton UP.
- Merkley, Eric, and Dominik Stecula. 2018. “Party Elites or Manufactured Doubt? The Informational Context of Climate Change Polarization.” *Science Communication* 40: 258-274.
- Montmarquet, James. 1993. *Epistemic Virtue and Doxastic Responsibility*. Lanham: Rowman and Littlefield.
- Nadarevic, Lena and André Aßfalg. 2017. “Unveiling the Truth: Warnings Reduce the Repetition-Based Truth Effect.” *Psychological Research* 81: 814-826.
- Nottelmann, Nikolaj. 2007. *Blameworthy Belief: A Study in Epistemic Deontology*. Dordrecht: Springer.
- Olsen, Kristian. 2017. “A Defense of the Objective/Subjective Moral Ought Distinction.” *The Journal of Ethics* 21: 351-373.
- Pantazi, Myrto, Mikhail Kissine, and Olivier Klein. 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, Rik. 2017. *Responsible Belief: A Theory in Ethics and Epistemology*. New York: Oxford University Press.
- Porter, Ethan, and Thomas Wood. 2019. *False Alarm: The Truth about Political Mistruths in the Trump Era*. Cambridge: Cambridge UP.
- Rapp, David. 2016. “The Consequences of Reading Inaccurate Information.” *Current Directions in Psychological Science* 25: 281-285.
- Rapp, David, Scott Hinze, Kristine Kohlhepp, and Rachel Ryskin. 2014. “Reducing Reliance on Inaccurate Information.” *Memory and Cognition* 42: 11-26.
- Richter, Tobias, Sascha Schroeder, and Britta Wöhrmann. 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.
- Rini, Regina. 2017. “Fake News and Partisan Epistemology.” *Kennedy Institute of Ethics Journal* 27: E-43-E-64.
- Sylvan, Kurt. 2020. “An Epistemic Nonconsequentialism.” *Philosophical Review* 129: 1-51.
- Tesler, Michael. 2018. “Elite Domination of Public Doubts About Climate Change (Not Evolution).” *Political Communication* 35: 306-326.
- Thorson, Emily. 2016. “Belief Echoes: The Persistent Effects of Corrected Misinformation.” *Political Communication* 33: 460-480.
- Unkelbach, Christian and Alex Koch. 2019. “Gullible but Functional? Information Repetition and the Formation of Beliefs.” In *The Social Psychology of Gullibility*, edited by Joseph Forgas and Roy Baumeister, 42-60. New York: Routledge.
- van der Linden, Sander, Anthony Leiserowitz, Seth Rosenthal, and Edward Maibach. 2017. “Inoculating the Public against Misinformation about Climate Change.” *Global Challenges* 1: 1600008.
- Verheggen, Bart, Bart Strengers, John Cook, Rob van Dorland, Kees Vringer, Jeroen Peters, Hans Visser, et al. 2014. “Scientists’ Views about Attribution of Global Warming.” *Environmental Science & Technology* 48: 8963-8971.
- Walter, Nathan, Jonathan Cohen, R. Lance Holbert, and Yasmin Morag. 2020. “Fact-

Checking: A Meta-Analysis of What Works and for Whom.” *Political Communication* 37: 350-375.

Whitcomb, Dennis, Heather Battaly, Jason Baehr, and Daniel Howard-Snyder. 2017.

“Intellectual Humility: Owning Our Limitations.” *Philosophy and Phenomenological Research* 94: 509-539.

Worsnip, Alex. 2019. “The Obligation to Diversify One’s Sources: Against Epistemic Partisanship in the Consumption of News Media.” In *Media Ethics, Free Speech, and the Requirements of Democracy*, edited by Carl Fox and Joe Saunders. New York: Routledge.

Zagzebski, Linda. 1996. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge: Cambridge UP.