



Epistemically blameworthy belief

Jessica Brown¹ 

Published online: 14 December 2019
© The Author(s) 2019

Abstract When subjects violate epistemic standards or norms, we sometimes judge them blameworthy rather than blameless. For instance, we might judge a subject blameworthy for dogmatically continuing to believe a claim even after receiving evidence which undermines it. Indeed, the idea that one may be blameworthy for belief is appealed to throughout the contemporary epistemic literature. In some cases, a subject seems blameworthy for believing as she does even though it seems *prima facie* implausible that she is morally blameworthy or professionally blame-worthy. Such cases raise the question of whether one can be blameworthy for a belief in a specifically epistemic sense rather than in some already recognised sense, such as being morally or professionally blameworthy. A number of authors have recently argued that there is a moral or social sense in which one ought to conform one’s beliefs to the evidence (e.g. Goldberg, Graham, Vanderheiden). In this paper, I argue that even while accepting that there are moral and social norms governing belief, there are cases in which a subject is blameworthy for a belief but isn’t plausibly morally or socially blameworthy. If this latter view is correct, then we may need to develop a new account of blame which can be applied to beliefs which are not morally or socially blameworthy.

Keywords Assertion · Belief · Epistemic norms · Blameworthy belief · Blame

✉ Jessica Brown
jab30@st-andrews.ac.uk

¹ Arché Philosophical Research Centre, University of St Andrews, St Andrews, UK

1 Introduction

When subjects violate epistemic standards or norms, we sometimes judge them blameworthy rather than blameless. For instance, we might judge a subject blameworthy for dogmatically continuing to believe a claim even after receiving evidence which undermines it. In some cases, the relevant blame seems to be moral or professional blame for the subject violates a moral or professional duty in believing against the evidence. For instance, a police detective might be both morally and professionally blameworthy for dogmatically continuing to believe that someone is responsible for a crime even after receiving evidence which undermines that claim. But, it's far from clear that all cases in which subjects are blameworthy for violating epistemic standards are cases in which they are morally blameworthy or professionally blameworthy. For instance, if a private citizen dogmatically clings to a belief, say, about the age of the earth, even after receiving evidence which undermines her belief, it is much less obvious that she is blameworthy in a moral or professional sense. For it is not obvious she has any moral or professional duty to conform her belief about the age of the earth to the evidence.

Such cases raise the question of whether one can be blameworthy for a belief in a specifically epistemic sense rather than in some already recognised sense, such as being morally or professionally blameworthy. This is an important question. On the one hand, the idea that one may be blameworthy for belief is appealed to throughout the contemporary epistemic literature and, on the face of it, the cases discussed don't always seem to involve moral or professional blame. On the other hand, a number of authors have recently argued that there is a moral or social sense in which one generally ought to conform one's beliefs to the evidence (e.g. Goldberg, Graham, Vanderheiden). If this latter view is correct, then existing accounts of blame may be applicable to blameworthy belief. But, if it is incorrect, then we may need to develop a new account of blame which can be applied to beliefs which are blameworthy in a distinctively epistemic sense. The aim of the paper is to evaluate whether there are cases in which one is blameworthy for a belief in a distinctively epistemic sense.

In Sect. 2, I start by providing several examples in which it's intuitive to think that a subject is blameworthy for failing to believe that *p* even though the evidence supports that *p*. The cases are set-up so that, unlike the example of the policeman above, the subject has no obvious moral or professional obligation to believe in accord with the evidence. In addition, it's stipulated that *p* is not relevant to any of her practical concerns. *Prima facie*, the cases are examples in which a subject is blameworthy in an epistemic sense, specifically that she is blameworthy for violating the norm that one epistemically ought to conform one's doxastic states to the evidence.

However, this initial case needs strengthening if it is to establish that subjects can be blameworthy in an epistemic sense. First, even those who accept the standard view that one epistemically ought to conform one's doxastic states to the evidence may argue that one ought to conform one's doxastic states to the evidence in a further sense, say a moral or prudential sense. More radically, some reject the

standard view and argue that the norm that one ought to conform one's doxastic states to the evidence is ultimately grounded in or derived from social or moral norms. For example, Goldberg claims that "epistemic norms themselves are grounded in and reflect the legitimate social (epistemic) expectations we have of one another" (2018: 2). On either of these views, someone might attempt to accommodate the cases by holding that the subject is blameworthy in an already established sense, such as being morally blameworthy. In the second half of the paper, I consider and reject a variety of such alternative explanations of the cases, whether they appeal to moral or social norms. I conclude that subjects may be blameworthy for their beliefs in a specifically epistemic sense.

2 Blameworthy and blameless belief

Throughout the paper, I will be assuming that one can be blameworthy for a belief to focus on the question of whether one can be blameworthy for a belief in a specifically epistemic sense—for violating a norm specifying what doxastic states one epistemically ought to have. Of course, some do deny that one can be blameworthy for belief. So it may be useful to quickly sketch a variety of different accounts that have been given of how one can be responsible and potentially blameworthy for belief. Scepticism about responsibility for belief is often motivated by the suggestion that belief fails to meet some necessary condition for responsibility, such as being under voluntary control or being something for which one is accountable to others. However, many reject these allegedly necessary conditions arguing that one can be responsible for things which are not under one's direct voluntary control such as forgetting or various emotional states (e.g. Smith 2005; Sher 2009); and/or denying that one is responsible for something only if one is accountable to others for doing it (e.g. Macnamara 2013; Smith 2015). More positively, a range of different accounts of how one is responsible for belief have been offered. Some defend responsibility for belief by appeal to the fact that beliefs are under our indirect control (e.g. Nettleman 2007; Peels 2017; Rettler 2018). Some defend answerability accounts according to which one is responsible for beliefs because they are the kind of states for which it makes sense to ask for a reason or justification (e.g. Scanlon 1998; Smith 2005; Hieronymi 2008). Others develop reasons-responsive accounts according to which one is responsible for beliefs because of one's ability to recognise and respond to reasons for belief (e.g. Ryan 2003; Steup 2008; Sher 2009; McCormick 2011; McHugh 2013).

From now on I will assume that one can be blameworthy for one's beliefs. I will be arguing that one can be blameworthy for one's beliefs in an epistemic sense, understood as being blameworthy for violating norms concerning what doxastic states one epistemically ought to have. More specifically, I will argue that one can be blameworthy for violating the norm that one ought to conform one's doxastic states to the evidence. It's standard within epistemology to suppose that this norm is an epistemic norm, underwritten by some epistemic value such as true belief or

knowledge.¹ Epistemic norms contrast with moral or prudential norms underwritten by moral or prudential value. The view that the norm of conforming one's doxastic states to the evidence is epistemic is motivated by cases with the following structure: the evidence makes it likely that *p* is true, but believing that *p* would be imprudent or immoral. Pascal's Wager illustrates the first possibility: Pascal argues that it is prudent to believe in God's existence to avoid the chance of suffering in hell even if one has no evidence of God's existence or even if the evidence supports that he does not exist. In another much discussed case, it's claimed that even though one has evidence that one's friend has done something wrong, reasons of loyalty make it the case that one morally ought not believe they've done wrong without further investigation. These cases suggest that we can distinguish a variety of senses in which one ought to have a belief. For instance in the Pascal example, from the perspective of prudence one ought to believe in God's existence. However, from the perspective of epistemic value, one ought not believe in God's existence for the fact that it's prudent to believe in his existence doesn't make it more likely to be true. Furthermore, one has evidence that he does not exist. Similarly, in the friendship example, that it would be disloyal to believe that one's friend has done wrong does not increase the likelihood that he hasn't done wrong; indeed one's evidence supports that one's friend has done wrong. *Prima facie*, these kinds of cases show that what one epistemically ought to believe can come apart from what one prudentially or morally ought to believe.²

Even if one accepts the standard view that one epistemically ought to conform one's doxastic states to the evidence, it's not straightforward to show that when a subject is blameworthy for failing to conform her doxastic states to the evidence she is blameworthy in an epistemic sense. For compatibly with that standard view, one could have a moral or professional duty to conform one's doxastic states to the evidence (as in the earlier example of the detective). As a result, from the facts that someone is blameworthy for believing against the evidence and that they epistemically ought to conform their doxastic states to the evidence, it doesn't follow that they are epistemically blameworthy for failing to do what they epistemically ought to do. For, they could be blameworthy for violating a moral or professional duty to conform doxastic states to the evidence. So to defend the claim that one can be blameworthy for failing to do what one epistemically ought to do, we need to find cases in which someone is blameworthy for failing to conform their doxastic states to the evidence but it doesn't seem that they have any moral or professional duty to conform their doxastic states to the evidence. I will now suggest two such cases.

First, consider Bonjour's (1980) case of Maud who in fact has a reliable clairvoyant power and uses that power to come to truly believe that the President is in New York. However, Maud also has evidence that the President is not in New

¹ Chignell, Andrew, "The Ethics of Belief", *The Stanford Encyclopedia of Philosophy* (Spring 2018 Edition), Edward N. Zalta (ed.), <https://plato.stanford.edu/archives/spr2018/entries/ethics-belief/>.

² Of course, some defend instrumentalist accounts of the norm of conforming one's beliefs to the evidence arguing that doing so generally promotes the satisfaction of one's desires and ends. For a detailed critique see Kelly (2003).

York (perhaps, the television is showing apparently live pictures of the President in Washington). Furthermore, Maud has been exposed to plenty of evidence that suggests that clairvoyance is impossible. Nonetheless, she dismisses the TV broadcast and the evidence against clairvoyance and continues to believe that the President is in New York. In so doing, it seems that Maud is blameworthy (e.g. see Williamson 2000, 2013, forthcoming; Sutton 2007; Hawthorne and Stanley 2008; Lasonen-Aarnio 2010; Hawthorne and Srinivasan 2013; Littlejohn forthcoming). However, at first sight, it doesn't seem that she's morally, professionally, or prudentially blameworthy for her failing. For unlike the example of the policeman considered earlier, it is not obvious that Maud has any moral or professional obligation to have true beliefs about the whereabouts of the President. Further, it's no part of the story that having a true belief about whether the President is in New York is important to any of Maud's practical projects or to the fulfilment of any of her goals or desires. Nevertheless, Maud remains intuitively blameworthy even if we fill the case out under the supposition that Maud had no particular moral or professional obligation to have true beliefs about the President's whereabouts, and even if having a true belief about the whereabouts of the President is irrelevant to Maud's practical projects.

It seems that a subject can be blameworthy not only for believing against the evidence but also suspending in the face of evidence. For instance, a subject, Trevor, might receive overwhelming evidence for some claim, consider it and yet suspend on the matter in question. For instance, he might receive overwhelming evidence that industrial pesticides used in agriculture are harming wild bees and yet suspend on whether they are harmful. Notice that it seems that Trevor is still blameworthy for suspending on the question of whether agricultural pesticides harm bees even if he has no particular moral or professional obligation to have true beliefs about this matter; and even if the issue is irrelevant to any of his practical projects.

To see that Trevor is blameworthy for suspending on the issue of the toxicity of pesticides to bees, it's important to distinguish suspending on an issue from mere failure to believe. We're constantly inundated with evidence about all matter of things. Given our cognitive limitations, it would be too demanding to expect that, in all such cases, we assess all this evidence and form appropriate beliefs. So, in some cases, there's nothing blameworthy about having evidence which strongly supports *p* yet forming no belief about *p* whatsoever. However, forming no belief about whether *p* is not to suspend on *p*. As Friedman (2013a) persuasively argues, we should reject "non-attitude" accounts of agnosticism or suspension as consisting in the lack of an attitude. For instance, she points out that it's not sufficient for being agnostic or suspending on *p* that one fails to have a belief either way. Cavemen were not agnostic about whether the result of the UK Brexit referendum would reduce the value of sterling even though they didn't have a belief about it either way. Nor is it sufficient to be agnostic about *p* that one has considered whether *p* and lacks a belief either way. For, one might have been interrupted before coming to any conclusion. Instead, Friedman argues that to suspend or be agnostic about *p* is to have an attitude that expresses or represents neutrality or indecision about which of *p* and not-*p* is true. Once we see that it is insufficient for suspending on *p* that one lacks both the belief that *p* and the belief that not-*p*, we see how Trevor can be blameworthy for

suspending in the face of overwhelming evidence even though we are often blameless for failing to believe one way or the other despite overwhelming evidence.

The examples of Maud and Trevor provide a good initial case that a subject can be epistemically blameworthy. More specifically, they seem to be blameworthy for violating the norm that one epistemically ought to conform one's doxastic states to the evidence. For they are intuitively blameworthy even though they don't have any obvious moral or professional obligation to conform the relevant doxastic states to the evidence; and the relevant doxastic states don't matter to any of their specific practical projects. Despite this, it may be suggested that the sense in which Maud and Trevor are blameworthy is prudential or moral. For it may be argued that in failing to have the belief which is supported by their evidence, they either potentially harm themselves later, or potentially mislead others. Let's consider these possibilities in order.

Although Maud and Trevor's failings don't directly impact any of their practical projects, it may be suggested that they have an indirect effect on their practical aims and are thus indirectly prudentially blameworthy. For instance, perhaps by believing that the President is in New York in the face of the evidence, Maud manifests or perhaps inculcates the bad epistemic habit of ignoring evidence. Such a habit can be expected to undermine the achievement of her practical goals in general. Thus, it may be claimed that the sense in which Maud is blameworthy for believing that the President is in New York against the evidence is prudential after all. A similar suggestion may be applied to the case of Trevor. However, this appeal to habit is problematic. It was no part of the description of either Maud or Trevor that they have the habit of ignoring evidence, or that the formation of the relevant doxastic state inculcates such a habit in them. Furthermore, it seems that they are blameworthy for the relevant doxastic states even if they don't manifest or inculcate a bad habit. The point is well known from the literature on morally blameworthy action: a subject can do something morally wrong, e.g. make a cruel remark, and be blameworthy for doing so, even without having the habit of making cruel remarks. As Sher (2006) argues, a good person can do bad things.³

Despite the failure of a habitual explanation of the sense in which Maud and Trevor are blameworthy, it might nonetheless be suggested that the sense in which they are blameworthy is immoral or prudential. For it may be suggested that by failing to have the beliefs that are supported by their evidence, they might harmfully mislead other agents directly or indirectly if others were to rely on them, or they might harm their own future doxastic states indirectly. However, we will see that problems face these proposed explanations.

Consider first the suggestion that by failing to have the beliefs that are supported by their evidence, Maud and Trevor might harmfully mislead other agents either directly or indirectly. For example, since Maud believes that the President is in New

³ Williamson (forthcoming) suggests a related strategy for judging subjects blameworthy for their beliefs by appeal not to whether the relevant belief exhibits bad epistemic habits but, instead, to whether the relevant belief is what someone who is in fact disposed to conform to the norms of belief would believe in the circumstances. I criticise this related account in Brown (2018a).

York, on being asked where he is, she might assert that he is in New York. Even though her assertion is true (since he is in fact in New York), this might potentially mislead someone into supposing that Maud has evidence that he is in New York. Since Trevor merely suspends on the issue of whether pesticides harm wild bees, he wouldn't assert that they do (or don't) harm wild bees on being asked. Instead, he might simply say that that he is not sure. In so doing, he might misleadingly suggest to others that his evidence on the matter is inconclusive.⁴ Thus it may be suggested that Maud and Trevor are blameworthy because in failing to have the beliefs that are supported by their evidence they risk misleading others.

However, the suggested explanation is problematic. First, at least to my mind, it misidentifies the reason why Maud and Trevor are blameworthy. Intuitively, they are blameworthy for failing to conform their doxastic states to the evidence and not because of the potential downstream effects of their doxastic states on others. Of course, not all will share this intuition. But, the suggestion also faces a second problem, namely that it overgeneralises. As we saw above, it's problematic to suggest that one is always blameworthy for failing to believe that *p* when the evidence strongly supports that *p*. We are constantly inundated with evidence about all kinds of subjects. Given our cognitive limitations, it is too demanding to suggest that we have a duty to assess all this evidence and form doxastic attitudes accordingly. So there need be nothing blameworthy about having evidence which strongly supports *p* and forming no belief either way about *p*. Instead, I argued above that one is blameworthy specifically for suspending on *p* when the evidence strongly supports *p*, where to suspend on *p* is not merely to fail to form a belief either way, but instead to have a doxastic attitude of neutrality on whether or not *p*. Once we recall that one is not always blameworthy for failing to form a belief about whether *p* when the evidence strongly supports *p*, we see that the proposed explanation overgeneralises. Suppose one fails to form any belief about *p* at a time *t*, even though one's evidence at *t* strongly supports that *p*. This might potentially result in one's misleading hearers if asked at a later time whether *p*. At the later time, one might have forgotten the evidence so that one cannot re-examine the issue. Since one has no belief about *p*, one might reply that one's not sure. But this might result in hearers falsely inferring that one's not sure because the evidence one had was inconclusive. On the proposed view that the possibility of misleading others makes one's doxastic state blameworthy, one will be blameworthy for failing to have the belief that *p*. Thus the proposed explanation of why Maud and Trevor are blameworthy overgeneralises.

A similar over-generalisation worry affects the suggestion that Maud and Trevor are blameworthy because their doxastic states might harm their future doxastic states indirectly. For example, it may be suggested that in suspending now on whether pesticides harm wild bees, Trevor may be led into later having the false belief that the evidence about the matter is inconclusive. After all, at the later time,

⁴ Notice that the suspension itself cannot be identified with metacognitive attitudes such as judging that the evidence is inconclusive on pain of ruling it out that unsophisticated cognisers (e.g. animals) can suspend (Friedman 2013b). As a result, on being asked whether *p*, someone who suspends on *p* may simply say that she is not sure rather than asserting that the evidence is inconclusive.

he may have forgotten the evidence he now has, so that he cannot assess it himself. Instead, realising that he suspends on the issue might lead him to suppose falsely that the evidence was inconclusive. In this way, his suspension might harm his future doxastic state. However, this potential explanation again overgeneralises. Suppose that an agent fails to form a belief about whether *p* at a time *t* even though at that time she has strong evidence that *p*. At a later time, realising that she has no belief about the matter but having forgotten the original evidence, she might erroneously conclude that she has no belief because the evidence was inconclusive at the earlier time. If the potential for such later errors makes the earlier failure to form a belief about whether *p* blameworthy, then this has the result that subjects are always blameworthy for failing to form a belief about whether *p* when they have strong evidence that *p*. But, as we saw earlier, that's not the case given our cognitive limitations.

We have now completed the initial case for the claim that subjects can be blameworthy in an epistemic sense for their beliefs. In particular, it's intuitive that Maud and Trevor are blameworthy for their doxastic attitudes and, in particular, for failing to conform their attitudes to the evidence. Furthermore, we have considered and rejected a range of potential explanations which seek to show that, given the details of the cases, they are morally, professionally or prudentially blameworthy for their doxastic attitudes. Despite this, the case in favour of epistemic blameworthiness is not yet complete. For, a range of authors have recently argued that there is a moral or social sense in which one ought generally to conform one's doxastic states to the evidence (Goldberg 2015, 2017, 2018; Graham 2015; Vanderheiden 2016). Indeed, Goldberg argues that the norm of conforming one's beliefs to the evidence is ultimately grounded in, or derived from, moral or social norms and is thus best regarded as moral or social. If there is a moral or social sense in which one ought generally to conform one's doxastic states to the evidence, then even absent any particular story according to which Maud and Trevor have moral, prudential or professional obligations to form the relevant doxastic states in accord with the evidence, there is still a moral or social sense in which they ought to do so. So to complete the argument that Maud and Trevor provide us with reason to suppose that one can be blameworthy in a specifically epistemic sense for failing to conform one's doxastic states to the evidence, we need to rule out these various alternative explanations for why they are blameworthy.

In the rest of the paper, I consider and reject various alternative explanations of why Maud and Trevor are blameworthy. I start in the next Sect. 3, with the idea that the norm of conforming one's doxastic states to the evidence is derived from our moral duties whether negative (not to harm others in general), or positive (a positive duty to inform others or assert only what the evidence supports). In Sect. 4, I discuss the suggestion that the norm of conforming one's doxastic states to the evidence is derived from, and grounded in, legitimate social epistemic expectations.

3 Deriving epistemic norms from moral norms

A number of philosophers have recently argued that one has a moral duty to conform one's beliefs to the strength of the evidence (e.g. Goldberg, Vanderheiden). The first way of defending this idea appeals to the claim that one can conform to one's moral duty not to harm others unnecessarily only if one is well-informed about whether one's actions are likely to harm others. For instance, from the moral duty not to harm a friend with a peanut allergy one might derive a moral duty to be well-informed about the ingredients of food one serves her, where this moral duty involves both acquiring evidence and forming beliefs in the light of that evidence. Initially, these familiar thoughts don't obviously seem to help us with the cases of Maud and Trevor since, as these cases were described, the subjects don't seem to have any moral duty to conform their doxastic states to the evidence. For instance, Maud doesn't seem to have any moral duty from which it follows that she morally ought to conform her belief about the whereabouts of the President to the evidence. However, some have recently argued that the moral duty not to harm others unnecessarily generates a general moral duty to be well-informed which has a wider scope than one might have imagined. As a result, it might be suggested that, despite initial appearances, Maud and Trevor do have a moral duty to conform their doxastic states to the evidence.

That we have a broad moral duty to be well-informed has been argued to follow from two features of the contemporary world: the global nature of international trade and politics and the easy availability of information via recent technologies such as the World Wide Web (Vanderheiden 2016). In particular, given the global nature of international trade and politics, one's duty to be well-informed extends much more widely than merely knowing about the local effects of one's actions. By buying clothes, I might be supporting child-labour in certain third world countries. By buying a car, I might be contributing toward climate change which negatively affects the globe. As a result, if I am to fulfil my moral duty not to harm others unnecessarily, I need to be well-informed about the prevalence of child-labour in the global garment industry and the effect of car-emissions on climate change. Furthermore, modern information technology including the World Wide Web, provides us with the means to acquire information about distant places and times, including the global garment industry and data on the world's climate.

However, even if we accept that these features of the contemporary world generate broad moral duties to be well-informed, it seems unlikely that these moral duties will extend to cover the full range of cases in which one ought to have doxastic states which conform to the evidence. Given our cognitive limitations, we are simply unable to be well-informed about everything. Nor do we end up with a cognitively manageable moral duty if we restrict it merely to having well-informed doxastic states where the evidence is easily available to us. As Vanderheiden puts it, "[i]n a world of abundant information, most small bits of knowledge are accessible to many persons at relatively low acquisition costs. But its totality exceeds anyone's cognitive capacities to effectively acquire and process, making some ignorance concerning facts relevant to harm prevention inevitable" (2016: 298).

It seems, then, that even if our duty to act morally generates a broad moral duty to be well-informed, this duty must somehow be restricted so that it does not universally apply. For instance, Vanderheiden's own suggested solution is that there are certain topics about which one has a moral obligation to inform oneself: "just as persons might be expected for prudential reasons to become more informed about threats to which they are particularly vulnerable, so also should they for moral reasons inform themselves about threats to others toward which they contribute as a kind of precautionary obligation" (305). But, if that's right, then there will be cases in which agents are blameworthy for dogmatically believing against the evidence but they don't have a moral duty to believe in accord with the evidence. For instance, consider Maud again who seems intuitively blameworthy for dogmatically believing that the President is in New York even in the face of evidence that he is in Washington and that clairvoyance is impossible. As we imagined the case, Maud's belief about the President doesn't contribute to any kind of threat whether to the President or anybody else. So Vanderheiden's suggested restriction would have the result that Maud has no moral duty to form beliefs about the President's whereabouts in conformity with the evidence. More generally, once we see that the moral duty to be well-informed must be constrained so that it doesn't universally apply, it won't explain all the kinds of cases in which subjects seem to be blameworthy for failing to conform their doxastic states to the evidence. Relatedly, one cannot derive the norm that one ought to conform one's doxastic states to the evidence from a moral duty not to harm others since the norm applies in cases in which there is no relevant derived moral duty.

Despite the failure of this first argument for the claim that one has a moral duty to conform one's doxastic states to the evidence, it is worth examining a second potential argument for this claim from the idea that one ought to help meet the needs of others, including their information needs.

3.1 A moral duty to meet others' information needs

In his book, *Assertion*, Goldberg appeals to the idea that we have a moral duty to help others with their information needs to derive a moral sense in which one ought to believe that *p* if and only if *p* is sufficiently supported by the evidence. First, Goldberg argues that there is a moral demand that we help others, including with their information needs. Suppose that it is mutually manifest to a speaker and hearer that the hearer needs information whether *p*, and this need is not defeated (call this "the assumption of undefeated information need"). Under this condition, Goldberg argues that the speaker morally ought to meet this need if she can. In particular, he argues that under these conditions, "S ought (morally) to address H's informational needs (regarding whether *p*) by asserting that *p* if and only if S occupies an epistemic position in which S can rule out all relevant alternatives to *p*" (197). Goldberg then argues that this moral requirement on assertion generates a moral requirement on belief. In particular, he argues that under the condition of undefeated information need, a speaker ought to believe that *p* if and only if she can rule out the relevant alternatives to *p*. First, given the tendency of speakers to assert what they believe, "if one is under conditions in which one's epistemic standing does not

warrant the assertion that *p*, then to believe that *p* is to unduly risk exposing others to the risk of unwarranted—and so unhelpful—assertion. Thus, one should not believe that *p* in these circumstances” (201). Second, he suggests that if one can rule out the relevant alternatives to *p* and so ought to assert that *p* to the hearer, then one ought to believe that *p* both because one has warrant to believe that *p* and because to assert without believing that *p* would be insincere (200). Thus, under the condition of undefeated information need, a subject ought to believe that *p* if and only if she can rule out the relevant alternatives to *p*.

As it stands, Goldberg’s argument establishes only a conditional moral requirement: that under the condition of undefeated information need, a subject ought to believe that *p* if and only if she can rule out the relevant alternatives to *p*. Thus, when this condition is not met, the relevant requirement on belief does not apply. As a result, even if there is a moral sense in which, in some circumstances, one ought believe that *p* if and only if one can rule out the relevant alternatives to *p*, this is hardly likely to extend to all beliefs. So, it fails to explain the sense in which, in many cases, a subject is blameworthy for not conforming her doxastic states to the evidence. To see this, reconsider clairvoyant Maud who continues to believe on the basis of her clairvoyant power that the President is in New York even despite the contrary evidence that the President is in fact in Washington and the evidence that clairvoyance is impossible. It seems that she is blameworthy for continuing to believe that the President is in New York in the circumstances even if it’s not mutual knowledge that any of her hearers need information on the whereabouts of the President. Thus, the sense in which knowledge of the hearer’s informational needs can generate a moral requirement to conform one’s beliefs to the evidence cannot explain the sense in which Maud’s belief about the President is blameworthy. Relatedly, one cannot hope to derive the norm that one ought to conform one’s doxastic states to the evidence from the moral duty to help others with their information needs since the norm applies even when there is no relevant information need.

Nonetheless, one might wonder whether reflecting on the way belief is linked to assertion could be used to defend an unconditional claim that one ought to conform one’s belief that *p* to the evidence irrespective of whether or not a hearer now needs information whether *p*. One way to try and do so would be to appeal to the fact that we often don’t know what particular informational needs others will have. As a result, it might be suggested that we have a general duty to enquire and gather information and store it in such a way that it can be used to answer the potential informational needs of others at a later time. For even if someone doesn’t now need it, it’s quite possible that someone will later need it. The suggestion, then, might be that we have a general moral duty to acquire good information. Thus, one ought morally actively enquire into the world. One ought morally believe that *p* if and only if one has strong evidence for *p*. One ought morally store those beliefs in case others later need that information.

However, even if there are moral norms governing assertion and belief derived from others’ informational needs, I will argue that they cannot provide an adequate account of epistemic blameworthiness. To see this, I start by considering our duty to meet others nutritional needs and then draw an analogy between that and our duty to

meet others' informational needs. If there is a general duty to help others it will surely include helping them meet their nutritional needs and not merely their informational needs. Of course, just as we cannot anticipate precisely what information others will need, we can't anticipate precisely what and how much food others will need. So if the difficulty of anticipating others' information need motivates acquiring and storing lots of varied information, so the difficulty of anticipating their nutritional needs would motivate acquiring and storing a wide variety and quantity of food. But, this seems far too demanding. Consider all the people currently hungry and in need of food in my town. Some need certain kinds of special food, such as gluten-free food, but it doesn't follow that I ought to store it just in case they need it. Even for those who don't need special kinds of food, it doesn't follow that I ought to have some food in the house just in case they stop by and ask for it. For, I could help meet their need even without having food stored in the house, say by enabling them to get food themselves (by giving them money) or telling them where they can get food (e.g. the supermarket or the local church food bank). Even if I do in fact have some food stored in the house, it doesn't follow that I should meet any fellow citizens' food needs. For some of my fellow citizens can easily feed themselves without any help from me. Even if some cannot feed themselves, such as children, it still doesn't follow that I ought to rush over and feed them. For, their parents have the primary duty to feed them. Others should step in only if their parents are not able or willing to do so. In general, then, even though it is plausible that we ought to help others, it's not plausible that we ought to acquire and store a vast variety of food just in case others need it.

Similar points apply to information need. Even if some of my fellow citizens need specialist information concerning cancer treatment, it doesn't follow that I, a non-doctor, ought to meet this need. I'm not best placed to provide this, and it's the responsibility of doctors in particular to provide such information to their patients. Even when I am well-placed to obtain some information which someone else needs, it doesn't follow that I ought to meet their need. For example, consider a student who wants to know the fine details of my University's appeal process. Even though as a University professor I'm well-placed to obtain this information, it doesn't follow that I should provide it to them. For, they may well be able to easily find it out for themselves by consulting the university's website. I may fulfil my moral and professional duties by telling them where they can find it rather than providing it to them myself. Just as in the case of food need, it's not generally true that I ought to obtain, store, and provide any and all information that my fellow citizens may potentially need. If there are duties to provide information, they are much more constrained than this.

In reply to these points, it may be suggested that there is a general sense in which we ought to meet the needs of others, but it is frequently defeated by other considerations such as the cost to ourselves, or the fact that those in need could easily meet the need themselves, or that others have a greater duty to meet the need. However, to the extent that the relevant ought is frequently defeated in this way, it doesn't provide a plausible account of blameworthy epistemic behaviour. For instance, take again the case of Maud who relies on her clairvoyant power in forming beliefs even in the face of contrary evidence and evidence against the

possibility of clairvoyance. According to the current suggestion, Maud's belief is blameworthy because she is failing to do what she ought to do to meet others' informational needs, in particular she fails to conform her beliefs to the evidence. But to the extent that we allow that the relevant ought is frequently defeated, then it may well be defeated in her case. For instance, perhaps she employs her clairvoyant power to form the belief that the President is in New York where, say, the Vice President needs to know where he is. Plausibly, others (say his security staff or his personal assistant) have a greater duty to meet the relevant information need, or perhaps the Vice President could easily meet it himself. Under such circumstances, the relevant ought would be defeated. But, Maud is still blameworthy in believing as she does. So, appeal to the general sense in which we ought to meet others' needs does not offer a plausible explanation of the sense in which Maud is blameworthy. Assuming that Maud is blameworthy in part because she violates norms governing belief, it seems we need to allow that there are norms governing belief other than the moral norm derived from the duty to help others with their information needs.

3.2 A moral duty to assert only what one's evidence supports

Even while admitting that we have no general moral reason to acquire, store and provide for the informational needs of others, it might be suggested that we have a more limited moral reason to believe only what is well supported by the evidence. Insofar as one morally ought to assert only what is supported by the evidence, it might be thought that this moral ought governing assertion generates a moral ought governing belief. To the extent that one is disposed to assert what one believes, if one's beliefs do not reflect the strength of one's evidence, then one's assertions won't either. Thus, in order to avoid asserting what is not supported by the evidence, one ought to ensure that one doesn't believe what is not supported by the evidence.

The proposed moral ought is far more constrained than those we have discussed already. It merely enjoins us not to believe on insufficient evidence. But it makes no claim that one ought to gather information, or that one ought positively to believe when one has strong evidence. It merely rules out believing on insufficient evidence. For although believing may dispose one to assert, failing to believe that *p* does not plausibly dispose one to assert either *p* or not-*p*. Similarly, failing to gather information whether *p* doesn't dispose one to form a belief concerning *p* that one is likely to assert. Relatedly, it doesn't rule out suspending on whether *p* when the evidence makes *p* certain, or almost certain. For, again, someone who suspends on whether *p* isn't thereby disposed to assert that *p* or that not *p*.

As a result of the constrained nature of the proposed moral ought, it fails to account for the full range of cases in which we judge subjects blameworthy for their epistemic behaviour. At best, it could serve as the ground of judgements we make that a belief is blameworthy when it is held in the face of the evidence. But, it does nothing to ground the idea that one ought not suspend on whether *p* when one has strong evidence that *p*. As a result, it cannot explain the sense in which certain cases of suspension seem blameworthy. Recall our earlier case of Trevor who acquires extremely strong evidence that agricultural pesticides are harming wild bees but nonetheless suspends on this question even despite his evidence. Intuitively, Trevor

is blameworthy for suspending on this question in the face of the evidence. However, it's difficult to explain this by appeal merely to the idea that one ought assert only that for which one has sufficient evidence and the fact that one is disposed to assert what one believes. For Trevor's suspension does not dispose him to assert that for which he lacks sufficient evidence. (Further I argued above against the suggestion that Trevor ought not suspend because his suspension might potentially result in him misleading others or his later self, say into believing or asserting that his evidence about pesticides and wild bees is inconclusive.)

We've now considered a variety of attempts to defend the idea that we morally ought to believe according to the strength of the evidence. However, we have seen that they fail to accommodate the full range of cases in which we ought to apportion our doxastic states to the strength of the evidence and may be blameworthy if we fail to do so. Relatedly, we have seen that we cannot derive the norm that one ought to conform one's doxastic states to the evidence from the relevant moral duties. Our discussion included appeal to the moral duty to avoid harming others unnecessarily, appeal to the moral duty to meet others' information needs, and appeal to the moral duty not to assert that for which one lacks sufficient evidence. In the next section, I consider whether one can derive the norm that one ought to conform one's beliefs to the evidence from social norms governing belief.

4 A social duty to meet epistemic expectations

A number of philosophers have recently defended the idea that there are social norms governing belief. For instance, Graham (2015) argues that there are social norms governing assertion and belief where he understands social norms as regularities which exist because prescribed. In particular, he argues that there is a social norm to provide true and relevant information, which generates a social norm to acquire evidence and believe on adequate evidence. For acquiring evidence and believing on adequate evidence is the way in which we can fulfil the social norm to provide true and relevant information. Goldberg (2015) argues that we are entitled to expect certain things of others, including that they meet certain moral and epistemic standards. Goldberg gives as examples of such epistemic expectations that others will form beliefs about readily observable matters of fact in their local environment and that others will form beliefs in epistemically responsible ways which he explains as including avoiding "the vices of closed mindedness, dogmatism, incuriosity and so forth". Furthermore, he argues that this is so even if one has no specific professional, institutional, or moral obligation to do so. Instead, he suggests that the sense in which one ought to, say, conform one's beliefs to the evidence is social or pragmatic (2891). More broadly, his project is to reverse what many regard as a natural way to think about the relation between epistemic norms and social expectations. While many hold that the epistemic norms governing belief explain or rationalise the expectations we have of one another as epistemic subjects, he instead defends the view that "epistemic norms themselves are grounded in and reflect the legitimate social (epistemic) expectations we have of one another" (Goldberg 2018: 2).

It's important to note that to say there is a social norm requiring something is not to say the social norm is legitimate. Relatedly, from the fact that there is a social norm requiring something it doesn't follow that one ought to do it. For social norms can be morally abhorrent such as norms favouring female circumcision. Goldberg agrees that it is only legitimate social expectations which are of interest and which potentially ground true claims that someone should have done something even if they failed to do it (Goldberg 2017: 2867). Goldberg attempts to show that the relevant epistemic expectations are legitimate by appeal to our practices of informational exchange and acting together with others. Summarising his argument that these expectations are legitimate, he says "such engagements are practically rational only if these expectations are legitimate; but these engagements *are* practically rational; so the expectations are legitimate" (2881). He attempts to show that the practical rationality of such engagements depends on the legitimacy of the expectations by arguing that if one calls into question the relevant expectations, one has called into question the relevant practice. For instance, he says "if you call into question e.g. that we are entitled to expect people to be epistemically responsible ... you have once again called into question something that is part and parcel of our everyday practice of relying on others say so: why should a hearer, H, take a speaker, S's, word for something if H cannot expect that S has avoided the vices of closed mindedness, dogmatism, incuriosity, and so forth?" (Goldberg 2017: 2894).⁵

Goldberg's argument hinges on the idea that if we call into question the claim that we are entitled to suppose that others will meet certain epistemic expectations, say, know the salient features of their immediate environment, this would undermine the practical rationality of the practice of testimony. But since he takes it that our practice of testimony is practically rational, he thinks it follows that we are entitled to the relevant epistemic expectation. However, we will see that his argument doesn't show that we are entitled to the relevant epistemic expectation in the required sense. To set up the problem, let's first pull apart two different notions of entitlement, an epistemic and a normative notion. For instance, suppose that one has a lazy colleague and a conscientious colleague with the same professional duties but a very different track record in terms of fulfilling them. It may be legitimate to expect them each to mark their scripts on time so that, if they fail to do so, it's true to say of each of them that they should have marked their scripts on time. However, given one's long experience with both of them, the evidence might support that only the conscientious but not the lazy colleague will mark their scripts on time. In the normative sense, one is entitled to expect both of them to mark their scripts on time. However, in the epistemic sense, one is entitled to expect only the conscientious colleague, but not the lazy colleague, to do so.

For the purposes of his argument, Goldberg needs to establish that the practical rationality of the practice of testimony shows that we are entitled to certain epistemic expectations in the normative, rather than the epistemic sense. Goldberg aims to appeal to the notion of entitlement to show how it can be true to say that a

⁵ For similar comments, see his discussion of testimony and acting with others in his 2018 book, 153–157.

subject should have known something, even when there is no epistemic reason to suppose that she does know it, or even if there is epistemic reason to suppose that she does not know it. For example, he wants to say that a known-to-be lazy doctor should know the best treatments for straightforward conditions commonly suffered by her patients even if, given her laziness, it is epistemically unlikely that she does. So, as Goldberg recognises, it is the normative sense of entitlement which is in play in his argument (2881).⁶ Further, it is the normative not the epistemic notion of entitlement which is potentially connected to criticism and blame of an agent for failing to meet our expectations.

Reflection on the practical rationality of testimonial exchanges can certainly be used to support the idea that we are epistemically entitled to suppose that others will meet our epistemic expectations. For instance, if we were to receive evidence that others have become infected with a disease which undermines the reliability of their basic belief-forming processes, such as vision, then this would undermine the practical rationality of relying on their testimony. Thus, one might argue that assuming it is practically rational to rely on testimony, we have epistemic entitlement to suppose that others will meet our epistemic expectations, say that their basic belief-forming processes are reliable. However, the fact that the practical rationality of our testimonial practices presupposes our having an epistemic entitlement to suppose that others will meet our epistemic expectations fails to show that we are normatively entitled to the relevant epistemic expectations.

To see this, consider the apocryphal story that Kant was so regular in taking a midday walk that his fellow citizens could tell the time by observing his walk. The practical rationality of relying on Kant to tell the time depends on his fellow citizens having epistemic entitlement to suppose that he takes his daily walk at midday. If his fellow citizens acquire evidence that he's become disabled through injury and unable to walk, then it's no longer practically rational to rely on him to tell the time. However, that the practical rationality of relying on Kant depends on their having epistemic entitlement to suppose that he walks daily at midday does not show that they are normatively entitled to suppose that he walks daily at midday. Indeed, even if we imagine that they lack any such normative entitlement, their practice of relying on Kant may still be practically rational.

For instance, imagine that Kant hasn't agreed to help his citizens by taking a walk at midday every day and that, caught up in his deep philosophical thoughts, he has no idea that fellow citizens have come to rely on him for telling when it is midday. In these circumstances, his fellow citizens cannot legitimately expect him to act as a timepiece for them. For instance, if one day he decides to take his walk at 2 pm rather than midday, there is no legitimate expectation which could ground the claim that he should have taken his walk at 12 pm, or to criticise him for failing to take his walk at 12 pm. Putting it into Goldberg's terms, in the circumstances, his fellow citizens don't have a normative entitlement to suppose that he will walk

⁶ He explains that "[t]he sort of expectation I have been (and will be) discussing at present is a normative one... That is why, if a subject were to fail to satisfy our expectation e.g. to be minimally observant of his surroundings, we would not take this failure to bear against our expectation that S be observant. Rather, we would say something like "In the future I expect more of you (better from you)" (2018:156).

every day at midday. Nonetheless, this needn't undermine the practical rationality of Kant's fellow citizens relying on his walk to tell the time so long as they are epistemically entitled to suppose that he takes his daily walk at midday, say because they know he has taken his daily walk at midday for the last 40 years and they have no evidence that he won't do so tomorrow.⁷ Thus, the fact that the practical rationality of relying on Kant requires that his fellow citizens are epistemically entitled to suppose that he walks daily at 12 pm fails to show that his fellow citizens are normatively entitled to suppose that he walks daily at 12 pm. Analogously, that the practical rationality of testimonial practices requires that we are epistemically entitled to suppose that others will meet our epistemic expectations fails to show that we are normatively entitled to suppose that others will meet our epistemic expectations. But we've seen that for Goldberg's argument to succeed, he needs to show that we are normatively entitled to suppose that others will meet our epistemic expectations.

We have now examined, and rejected, a number of arguments that might be taken to show that the sense in which a subject is blameworthy for believing against the evidence is moral, social, professional, or practical. Some suggest that one has a moral duty to believe on adequate evidence derived from some independent moral duty. Suggestions for the relevant moral duty include one's duty not to harm others unnecessarily; one's duty to help others with their information needs; and moral duties regarding assertion. All these suggestions suffer from the similar problem that even if there are such moral duties, they do not plausibly extend to cover the full range of cases in which we judge subjects blameworthy for failing to have well-informed beliefs. Relatedly, one cannot hope to derive the norm that one ought to conform one's doxastic states to the evidence from these moral duties since the norm applies in some cases in which one has no such moral duty. Some instead suggest that there is a social or pragmatic sense in which one ought to conform one's beliefs to the evidence. In particular, Goldberg attempts to derive such a social or pragmatic ought from our practice of testimony and argues that the epistemic norm that one ought to conform one's doxastic states to the evidence derives from our legitimate social epistemic expectations. But I have argued that the derivation fails. The practice of testimony shows at best that there is a general epistemic entitlement to suppose that others are epistemically responsible in the management of their beliefs. However, it fails to show that we are normatively entitled to suppose that others are epistemically responsible. But it is only a normative notion of entitlement which would ground criticism of an agent for dogmatically believing against the evidence in circumstances in which she's epistemically unlikely to meet the relevant requirement, or in which one knows that she hasn't met it. I conclude that these various arguments fail to show that whenever a subject is blameworthy for

⁷ Notice that is so even if Kant's fellow citizens recognise that in the circumstances they are not normatively entitled to suppose that Kant will walk every day at 12 pm. E.g. even if they recognise that Kant hasn't agreed to act as a timepiece for them and is unaware that they are relying on him, their practice of relying on him may still be practically rational to the extent that they have epistemic entitlement to suppose that he walks daily at midday.

failing to conform her doxastic states to the evidence she is blameworthy for violating a moral, social, professional, or practical norm.

5 Conclusion

The idea that a subject may be blameworthy for her beliefs is widely appealed to throughout a range of contemporary epistemic debates concerning justification, knowledge, defeat, higher-order evidence, akrasia and peer disagreement. Here we've been examining in what sense, or senses, a subject may be blameworthy for her beliefs. In particular, we've been examining a variety of attempts to argue that the sense in which one is blameworthy for failing to conform one's doxastic states to the evidence can be understood by appeal to the idea that one is blameworthy for violating a moral, social, professional or practical norm. We have examined a variety of attempts to derive the norm that one ought to conform one's beliefs to the evidence from independent moral norms. These include the norms that one ought not harm others unnecessarily; that one ought to help others with their information needs; and that one ought to assert only what is supported by adequate evidence. In addition, we considered the ideas that, in a social sense, one ought to believe according to the evidence. While not denying that there may be such moral or social oughts, we saw that they do not cover the full range of cases in which we judge subjects blameworthy for belief. I suggest, then, that a subject can be blameworthy in a specifically epistemic sense, for violating an epistemic norm specifying that she epistemically ought to conform her doxastic states to the evidence.

If that is right, then this leaves us with the project of understanding the epistemic sense in which subjects may be blameworthy and blamed for their beliefs. There is a large literature in ethics on what it is to be morally blameworthy for one's actions. Relatedly, there is a large literature on what it is to morally blame someone. However, there is little developed work on what it is to epistemically blame someone for their beliefs.⁸ Moreover, some of the key features in terms of which moral blame is standardly understood don't seem applicable to all cases of epistemically blameworthy belief. A standard element of popular "conative" accounts of moral blame is a judgement or belief that the subject acted in a way which is morally wrong (e.g. Scanlon, Sher). But as we've seen, there are plausibly cases of blameworthy belief that don't involve moral wrongdoing. Further, moral blame is typically explicated in terms of sanction, resentment, guilt, apology and compensation. But, we don't punish people for their beliefs and it would seem inappropriate for Maud to feel guilty about her dogmatic belief, to apologise or compensate others, or for others to resent her for her belief. So, the conclusion that

⁸ There are some exceptions: Peels (2013) argues that doxastic blame is an affective attitude which involves holding others responsible for their beliefs; Rettler (2018) argues that doxastic blame should be understood as involving the demand that someone acknowledge their failure to conform to epistemic norms; Brown (2018b) provides an account of epistemic blame as consisting in a characteristic set of emotional and behavioural dispositions unified by a certain belief-desire pair. Kauppinen (2018) defends a notion of epistemic accountability for violating epistemic norms which he argues is distinct from blame.

we can be blameworthy in a specifically epistemic sense for belief leaves us with the new project of trying to understand what it is to blame someone for such a belief.

Acknowledgements Thanks for a Mind Association Senior Research Fellowship (2017–2018) which enabled me to finish this paper. I presented earlier versions of this material at a variety of events and received helpful comments from participants, including talks in Edinburgh and Hamburg; at the 2017 St Andrews Arché workshop on blame, the 2017 Bled conference, and the 2019 Urbino Epistemology Summer School. Thanks also to Sandy Goldberg, Peter Graham and Matt McGrath for helpful discussion and feedback.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Bonjour, L. (1980). Externalist theories of empirical knowledge. *Midwest Studies in Philosophy*, 5(1), 53–74.
- Brown, J. (2018a). *Fallibilism: Evidence and knowledge*. Oxford: Oxford University Press.
- Brown, J. (2018b). What is epistemic blame? *Nous*. <https://doi.org/10.1111/nous.12270>.
- Dennett, D. C. (1986). Information, technology, and the virtues of ignorance. *Daedalus*, 115(3), 135–153.
- Friedman, J. (2013a). Suspended judgement. *Philosophical Studies*, 162(2), 165–181.
- Friedman, J. (2013b). Question-directed attitudes. *Philosophical Perspectives*, 27, 145–174.
- Goldberg, S. C. (2015). *Assertion: On the philosophical significance of assertoric speech*. Oxford: Oxford University Press.
- Goldberg, S. C. (2017). Should have known. *Synthese*, 194(8), 2863–2894.
- Goldberg, S. C. (2018). *To the best of our knowledge*. Oxford: Oxford University Press.
- Graham, P. J. (2015). Epistemic normativity and social norms. In D. K. Henderson & J. Greco (Eds.), *Epistemic evaluations*. Oxford: Oxford University Press.
- Hawthorne, J., & Srinivasan, A. (2013). Disagreement without transparency: Some bleak thoughts. In D. Christensen & J. Lackey (Eds.), *The epistemology of disagreement* (pp. 9–30). Oxford: OUP.
- Hawthorne, J., & Stanley, J. (2008). Knowledge and action. *Journal of Philosophy*, 105, 571–590.
- Hieronymi, P. (2008). Responsibility for believing. *Synthese*, 161(3), 357–373.
- Kauppinen, A. (2018). Epistemic norms and epistemic accountability. *Philosophers' Imprint*, 18(8), 1–16.
- Kelly, T. (2003). Epistemic rationality is instrumental rationality: A critique. *Philosophy and Phenomenological Research* LXVI, 3, 612–640.
- Lasonen-Aarnio, M. (2010). Unreasonable knowledge. *Philosophical Perspectives*, 24, 1–21.
- Littlejohn, C. (Forthcoming). A plea for epistemic excuses. In F. Dorsch & J. Dutant (Eds.), *The new evil demon*. OUP: Oxford.
- Macnamara, C. (2013). Taking demands out of blame. In J. Coates & N. A. Tognazzini (Eds.), *Blame: Its nature and norms* (pp. 141–161). New York: Oxford University Press.
- McCormick, M. (2011). Taking control of belief. *Philosophical Explorations*, 14(2), 169–183.
- McHugh, C. (2013). Epistemic responsibility and doxastic agency. *Philosophical Issues*, 23, 132–157.
- Nettleman, N. (2007). *Blameworthy belief: A study in epistemic deontologism*. Dordrecht: Springer.
- Peels, R. (2013). Against doxastic compatibilism. *Philosophy and Phenomenological Research*, 89(3), 679–702.
- Peels, R. (2017). *Responsible belief*. Oxford: OUP.

- Rettler, L. (2018). In defence of doxastic blame. *Synthese*, 195(5), 2205–2226.
- Ryan, S. (2003). Doxastic control and the ethics of belief. *Philosophical Studies*, 114(1–2), 47–79.
- Scanlon, T. M. (1998). *What we owe to each other*. Cambridge, MA: Harvard University Press.
- Sher, G. (2006). *In praise of blame*. Oxford: OUP.
- Sher, G. (2009). *Who knew? Responsibility without awareness*. Oxford: OUP.
- Smith, A. M. (2005). Responsibility for attitudes: Activity and passivity in mental life. *Ethics*, 115, 236–271.
- Smith, A. M. (2015). Responsibility as answerability. *Inquiry*, 58(2), 99–126.
- Steup, M. (2008). Doxastic freedom. *Synthese*, 161, 375–392.
- Sutton, J. (2007). *Without justification*. Cambridge, MA: MIT Press.
- Vanderheiden, S. (2016). The obligation to know: Information and the burdens of citizenship. *Ethical Theory and Moral Practice*, 19, 297–311.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: OUP.
- Williamson, T. (2013). Response to Cohen, Comesana, Goodman, Nagel and Weatherson on Gettier cases in epistemic logic. *Inquiry*, 56, 77–96.
- Williamson, T. (Forthcoming). Justifications, excuses, and sceptical scenarios. In F. Dorsch & J. Dutant.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.