



What is epistemic blame?

Jessica Brown

Date of deposit	05 11 2018
Document version	Author's accepted manuscript
Access rights	Copyright © 2018 Wiley Periodicals, Inc. This work is made available online in accordance with the publisher's policies. This is the author created, accepted version manuscript following peer review and may differ slightly from the final published version.
Citation for published version	Brown, J. (2018). What is epistemic blame? <i>Noûs</i> , <i>Early View</i> .
Link to published version	https://doi.org/10.1111/nous.12270

Full metadata for this item is available in St Andrews Research Repository at: <https://research-repository.st-andrews.ac.uk/>

WHAT IS EPISTEMIC BLAME?

Jessica Brown.

Arché Philosophical Research Centre, University of St Andrews.

To appear in Nous.

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. Indeed, it has become standard to appeal to the idea that subjects can be blameworthy or blameless for their beliefs in a variety of recent debates in epistemology including debates concerning the norm for belief, higher-order evidence, peer disagreement, epistemic akrasia, the new evil demon problem and defeat¹. One challenge facing such appeals is to provide an account of the conditions under which subjects are responsible and potentially blameworthy for their beliefs. A number of philosophers have attempted to provide just such an account (e.g. Smith 2005, Nottlemann 2007, Hieronymi 2008, McCormick 2011, McHugh 2013, and Peels 2017). By contrast, my focus is on a distinct challenge, namely to provide an account of what it is to blame someone for their beliefs. After all, if a subject is blameworthy then it is fitting to blame her. While there is a voluminous literature on blame, it mostly focuses on moral blame.² But when a subject is blameworthy for her beliefs, it's far from clear that she is blameworthy in the moral sense. Relatedly, when we blame a subject for her beliefs, it's far

¹ For example, see Williamson 2000, 2013, 2014, forthcoming; Hawthorne and Stanley 2008; Hawthorne and Srinivasan 2013; Lasonen-Aarnio 2014 and forthcoming; Srinivasan 2015; Littlejohn forthcoming; Weatherson forthcoming.

² For example, see the papers in Coates and Tognazzini 2013.

from clear that we morally blame her. The focus of this paper is an attempt to provide an account of what it is to blame someone for her beliefs and how it relates to moral blame.

In section 2, I set the stage for the discussion by providing an example in which a subject is intuitively blameworthy for a belief, where the sense in which she is intuitively blameworthy is not plausibly moral but epistemic. I sketch the way in which other authors have already argued that one can be responsible and potentially blameworthy for one's beliefs. I also outline the general approach to blame which I find most congenial—namely a variable account of blame on which it can take a wide variety of forms and needn't always take the form of sanction, resentment and guilt. Such variable accounts help make space for the notion of blame for belief. In section 3, I explain how a range of different approaches in epistemology have a use for the notion of blameless/blameworthy belief in addition to such traditional distinctions as a belief's being justified/unjustified, known/unknown, or in accord with the norm for belief or not. In section 4-6, I develop a new account of epistemic blame inspired by Sher's account of moral blame. Sher's account focuses on blame for morally bad action and morally bad character. But, belief doesn't appear to be either an action or a character trait. Furthermore, the cases of blameworthy belief on which I will focus don't seem to be cases of moral wrong nor do they seem to involve violating a moral norm. As a result, Sher's original account of blame isn't directly suitable for the relevant cases of blameworthy belief. Instead, I provide a new account of blame which shares important general features with Sher's account, but differs in the content of the belief-desire pair and the characteristic dispositions in terms of which blame is understood.

2. Blameworthy belief.

Other authors have recently defended the idea that we can be blameworthy for our beliefs (e.g. Smith 2005, Nettleman 2007, Hieronymi 2008, McCormick 2011, McHugh 2013, and Peels

2017). Since my focus is to develop an account of what it is to blame someone for her beliefs, I will not repeat their arguments in detail, but merely sketch the kind of frameworks they employ. In particular, those who defend the idea that we can be blameworthy for our beliefs reject the idea that being responsible or blameworthy for X requires that X is under one's voluntary control and one is accountable to others for doing X.³ More positively, those who accept that we can be blameworthy for our beliefs have adopted a range of different accounts of responsibility on which responsibility doesn't require voluntary control or accountability to others. These include attributability and answerability accounts.⁴

In addition to holding with others that we can be blameworthy for beliefs, I also endorse a variable account of blame on which blame need not take the form of sanction or the particular negative reactive attitudes of resentment and guilt (see e.g. Scanlon 1998 and 2008, Hieronymi 2004, Sher 2006, McGeer 2013, and Smith 2013). While it's true that we sometimes punish and rebuke those judged morally blameworthy, at other times we don't. For instance, I might blame my boss for his treatment of me, but judge it imprudent to punish and rebuke him. Similarly, blame needn't always take the form of resentment and guilt. For instance, I might blame a friend who's betrayed me by deciding not to continue the friendship with feelings of loss and disappointment rather than indignation and resentment. As well as

³ Thus, they reject accounts of responsibility as accountability provided by, inter alia, Wallace 1994 and 2011, Watson 1996, Levy 2005, and Darwall 2006.

⁴ Watson (1996) famously distinguished responsibility in the accountability and attributability senses. Answerability accounts of responsibility for belief have been offered by Scanlon 1998 and 2008, Smith 2005 and 2015, and Hieronymi 2008 and 2014. Other authors place different conditions on responsibility for belief such as reasons-responsiveness, e.g. McCormick 2011 and McHugh 2013; or indirect control, e.g. Nettleman 2007 and Peels 2017.

allowing blame to take various forms, two further important criteria for an account of blame are that 1) it allows for private blame (e.g. I might hide my blame out of fear); and 2) allows that blame goes beyond judging blameworthy. For one can judge that a subject is blameworthy and yet find oneself psychologically unable to blame her, or fail to blame her because one lacks the standing to do so (perhaps one is guilty of just the same misdemeanour so that blaming her would be hypocritical).

Variable accounts of blame are both plausible and also help make room for the idea of blaming people for their beliefs. We don't typically punish people for their beliefs. Further, some of the reactive attitudes often associated with blame, such as resentment and guilt, don't always seem appropriate in the case of blameworthy belief. To see this, consider a familiar case of blameworthy belief from the literature, Bonjour's case of Maud. Maud has a reliable clairvoyant power and uses that power to come to believe that the President is in New York. However, she also has strong evidence that the President is not in New York (perhaps, the television is showing apparently live pictures of the President in Washington). Furthermore, Maud has no evidence for the existence of her clairvoyant power, and has plenty of evidence that supports that clairvoyance is impossible. Importantly, Maud's problem is not that she simply cannot appreciate the strength of the contrary evidence. Like the rest of us, she understands that the television's showing apparently live pictures of the President in Washington is excellent evidence that he is not in New York. Furthermore, she understands the scientific evidence which supports that clairvoyance is impossible. Nonetheless, she dismisses the TV broadcast and the evidence against clairvoyance and dogmatically continues to believe that the President is in New York. In so doing, it seems that Maud is blameworthy. Consequently, we may well blame her for her belief. For instance, her close friends may rebuke her saying, "How can you be so stupid as to believe the President's in New York when the television is showing that he is in Washington?", or "How can you

believe the President is in New York? You know the evidence about clairvoyance”. When Maud attempts to explain away the compelling evidence against her belief, they may get angry with her and her conspiracy theories.

Notice that even if we rebuke and get angry with Maud, we wouldn’t punish her for her belief. Furthermore, it would seem inappropriate to resent her for her belief or for her to feel guilty about it. Nor would we demand that she compensate us or apologise for believing as she does. However, once we endorse a variable account of blame, the fact that we wouldn’t punish her, resent her or demand compensation/apology doesn’t show that our reaction to her is not a case of blame.⁵

While variable approaches to blame help make space for blame for belief, we will see that existing accounts of blame need to be extended to accommodate blameworthy belief. Existing accounts have been mainly developed with morally blameworthy action in mind. But, *prima facie*, belief is not an action.⁶ Furthermore, cases of blameworthy belief such as

⁵ There may be other cases in which blameworthy belief does lead to resentment and demands for apology. E.g. if a close friend or lover believes on insufficient evidence that I’ve had an affair or deceived them, I might be resentful and ask for an apology. However Maud’s case illustrates that not all cases of blameworthy belief seem to involve this.

⁶ While most deny that beliefs are actions, one prominent view—Sosa’s—treats a certain kind of belief—“judgemental belief”—as an instance of action. However, even on this view, it is problematic to simply apply existing accounts of blame to blameworthy beliefs. For a belief may be blameworthy even if it doesn’t seem to violate any moral norm but existing accounts of blame focus on specifically moral wrongdoing.

Maud's don't intuitively seem to involve moral wrongdoing.⁷ It's far from clear that Maud has a moral responsibility to conform her beliefs about the President's whereabouts to the evidence. Indeed, Clifford's view that there is a moral responsibility to believe in accord with the evidence is rejected by most contemporary epistemologists⁸. Nor need we suppose that Maud has any professional responsibility to form well-supported beliefs about the President's whereabouts. So, it seems implausible that we can explain the sense in which she is blameworthy by appeal to violation of moral or professional norms. I will use the expression "epistemic blame" to describe our reaction to Maud's dogmatic belief. Later I will develop an account of epistemic blame by extending existing accounts of moral blame. However, before doing so (sections 4-6), in the next section I will sketch why the idea of blameless/blameworthy belief should be attractive to a range of different approaches in contemporary epistemology.

3. Blameless/blameworthy belief and epistemology.

Even if we are convinced that one can be responsible for one's beliefs, and that we should negatively evaluate Maud's belief, we might reasonably wonder whether we can accommodate everything we want to say about Maud without appealing to the notion of blameworthy belief. In particular, we might hope to do so by appealing to some already-established notion in epistemology, such as unjustified belief or belief that violates the norms

⁷ Of course, some examples of blameworthy belief do have a moral dimension. For instance, we may judge a developer who believes against the evidence that certain tower block cladding is safe, leading to a disastrous tower block fire, morally blameworthy for their belief. But, it's far from clear that Maud is morally blameworthy for her belief.

⁸ Though see Goldberg 2015 and forthcoming for a more sympathetic view.

of belief. However, in this section, I will argue that a wide range of different approaches in contemporary epistemology have use for a notion of blameless/blameworthy belief in addition to such already-established notions.

To see this, consider the popular idea that truth is the norm of belief. A subject could violate this norm by falsely believing that *p* in such circumstances that we wouldn't judge her negatively, for example if her evidence strongly supports *p* ("case 1"). Alternatively, a subject could violate the truth norm by believing that *p* falsely in such circumstances that we would judge her negatively, for example if she believes that *p* dogmatically against the evidence ("case 2"). Since both cases involve the subject violating the truth norm yet in only the second is the subject blameworthy, we cannot identify being blameworthy with violating the norm of belief. An alternative way of dealing with our different reactions to the cases would be to suggest that, in case 1, the subject's belief is justified whereas, in case 2, the subject's belief is unjustified. So it might be suggested that a belief is blameworthy if and only if it is unjustified.

However, some radical contemporary externalists reject the traditional notion of justification on which one can have justification for false beliefs.⁹ As a result, they cannot explain why we negatively evaluate the subject in the second case but not the first case by claiming that, in the second case the subject's belief is unjustified, but in the first it is

⁹ Some of those who reject the idea that one can have justified false beliefs hold that a belief is justified if and only if it meets the norm for belief which they take to be knowledge (Sutton 2007; Williamson 2014 and forthcoming; Littlejohn forthcoming). Nonetheless, they suggest that one can conform to the knowledge norm and be blameworthy, or violate it and be blameless. So, they cannot explain being blameless/blameworthy by appeal to the notion of being justified/unjustified or norm-conformity/violation.

justified. For by stipulation, in the first case the subject's belief is false. In order to reflect our evaluations of agents, these radical externalists appeal to a notion of a belief's being blameless/blameworthy where this is distinct from its being justified/unjustified or conforming to/violating the norm for belief. These authors appeal to the notion of blameless/blameworthy belief to defend a wide range of externalist views including externalist norms of belief (such as a truth or knowledge norm)¹⁰; and externalist accounts of doxastic justification as knowledge¹¹.

The notion of blameless/blameworthy belief is not only useful to such hyper-externalists but also more traditional externalists who accept the notion of false justified belief. Externalists traditionally accept a reliabilist account of justification. But, in principle, there seems no reason to suppose that whether or not a belief was formed by a reliable process will align with our intuitions about whether or not the subject was blameless in forming that belief. This is even more obviously the case for those contemporary externalists who adopt a simple reliabilist account of justification without adding any anti-defeat condition. As a result, they hold that believing that *p* on the basis of a reliable process—say perception—is sufficient for the belief that *p* to be justified, even if the subject also has evidence against *p*, or that the process by which she formed the belief that *p* is unreliable. But intuitively a subject who continues to believe that *p* in these circumstances would be negatively evaluated. So, such externalists need a notion of blameless/blameworthy belief which cannot be identified with the distinction between justified/unjustified beliefs. As a result, in a case in which a subject believes on the basis of what is in fact a reliable process

¹⁰ E.g. Williamson 2000, Hawthorne and Stanley 2008, and Hawthorne and Srinivasan 2013.

¹¹ Williamson 2013, forthcoming; Littlejohn forthcoming.

but has putative defeating information, these externalists are motivated to claim that the belief is justified but nonetheless blameworthy.¹²

The notion of blameless/blameworthy belief is not only helpful to such externalist views, but should also be attractive to internalists. For instance, consider a more internalist friendly norm for belief such as the norm that one ought to conform one's beliefs to the evidence. A subject could violate this norm but nonetheless be intuitively blameless. For instance, consider a subject who falsely believes that *p* as a result of attractive but fallacious reasoning and who lacks the ability to see for herself that it is fallacious and has never been taught any different. An internalist could accept that the subject's belief that *p* is unjustified but nonetheless regard the subject as blameless for believing that *p*. By contrast, another subject who falsely believes that *p* as a result of the same fallacious reasoning might be blameworthy for doing so if she has been taught to avoid such fallacious reasoning and has the ability to identify the reasoning as fallacious. Thus, we can imagine cases of false belief via fallacious reasoning which are blameworthy and cases which are blameless. In that case, the internalist needs a notion of blameless/blameworthy belief which is distinct from the notion of justified/unjustified belief or norm conforming/violating belief.¹³ For both of our imagined subjects who believe that *p* as a result of fallacious reasoning end up with unjustified beliefs which violate the evidence norm of belief. But we judge them differently.¹⁴

¹² Anti-defeat views have been defended in Greco 2010; Lasonen-Aarnio 2010, forthcoming; Hawthorne and Srinivasan 2013; Baker-Hytech and Benton 2015; and Beddor 2015.

¹³ For instance, Owens (2000) accepts that a belief can be unjustified yet blameless in the course of defending his internalist approach.

¹⁴ The role of ability in determining whether we judge a subject blameworthy for, say, believing against the evidence undermines the recent suggestion made by some externalists

I've now argued that epistemologists endorsing both externalist and internalist approaches find useful a notion of blameless/blameworthy belief in addition to such traditional notions as justified/unjustified belief or norm-conforming/violating belief. Even while accepting that, an epistemologist might attempt to defend a deflationary view on which the relevant subjects are not genuinely blameworthy or blameless for their beliefs. Instead, it might be suggested that our reaction to these believers can instead be fully captured by appeal to instrumental considerations to do with attempting to improve subjects' conformity to the relevant epistemic norms.

In reply, it seems that this instrumental approach is both ill-motivated and unattractive. It would be motivated if there was some reason to think that beliefs fail to meet some necessary condition for responsibility. But, I started by sketching out a range of views on which beliefs do meet the conditions for responsibility. Further, I will argue that this instrumental approach is unattractive in the case of belief for reasons similar to those which make it unattractive in the case of action. In the case of action, treating someone in a way which improves the likelihood of them conforming to moral norms in the future needn't seem like a case of blame at all. For instance, medical treatment with drugs or brain surgery which improves norm conformity does not seem like blame. Or consider the kind of educative

that a belief is blameworthy if it manifests some bad epistemic habit or strategy (e.g. Lasonen-Aarnio 2010 and Hawthorne and Srinivasan 2013). For the beliefs of both fallacious reasoners exhibit bad epistemic habits or strategies. However, it would surely be preferable if these hyper-externalists adopt an account of blameworthy belief on which the degree of blame can be affected by one's ability to recognise and avoid bad reasoning. Further, doing so doesn't undermine their employment of the notion of blameworthy belief to deal with victims of evil demons or Gettier cases.

dispassionate conversation one might have with a small child in which one attempts to get them to consider what they have done, how their action made the other child feel, and finally to draw them into a moral assessment of what they've done. This dispassionate conversation doesn't seem like blame even though presumably we engage in it to improve norm conformity. Furthermore, blaming someone by sanctioning them, rebuking them and getting angry with them can sometimes be inappropriate or unfair even if to do so would be instrumentally effective. For instance, finding and blaming a scapegoat for something is unfair even if instrumentally effective. Likewise, it is surely plausible that different responses are appropriate to the wrongdoing of children and adults even if reacting to children in the same way as adults were to be instrumentally effective. For instance, we don't give children the same kinds of legal sentence that we give adults and it would seem unfair to do so, even if doing so turned out to be instrumentally effective.

For similar reasons, a merely instrumental account of our reactions to blameworthy believers is inadequate. Our negative reaction to blameworthy believers goes beyond a merely instrumental response. To see this, it is useful to compare our two subjects who each fall for fallacious reasoning. One of the subjects has the intellectual ability to see for herself that the reasoning is fallacious and moreover has been trained to avoid it. By contrast, the other subject does not have the intellectual ability to see for herself that the relevant fallacy is a fallacy, and has not been taught that the reasoning is fallacious. It would be inappropriate and unfair to rebuke and be angry with the second low-ability subject who cannot see for herself that it is fallacious reasoning, and has not been taught any different. That's so even if angrily rebuking the low-ability subject is instrumentally effective, say because it diminishes the low-ability subject's confidence so much that she avoids anything but the simplest reasoning and/or always relies on others for more complex reasoning. By contrast, it seems appropriate and fair to be angry with and rebuke the first high-ability subject for her failure to

avoid the fallacious reasoning given her intelligence and training. Instead of angrily rebuking the low-ability subject, we might offer to explain to her where her reasoning went wrong and offer her techniques for avoiding such fallacious reasoning in the future. Calmly offering such an explanation and training does not amount to blame even though it may help her conform to good norms of reasoning in the future. A second example which shows the difference between blaming someone for her belief and a merely instrumental response involves scapegoating. Just as scapegoating someone for a crime she didn't commit is unfair even if instrumentally effective, so is angrily rebuking a subject for committing a simple fallacy which she didn't in fact commit. Doing so is unfair even if instrumentally effective (imagine scapegoating a professional, such as an engineer or doctor, for committing a fallacy she didn't in fact commit to encourage others to avoid it). In sum, our reactions to blameworthy belief are not well captured by supposing they consist in a purely instrumental response.

It's worth noting that our different reaction to the low and high-ability subjects does place some constraints on the best account of how we are responsible for our beliefs. In particular, it may undermine simple versions of answerability accounts on which one is responsible for a belief because it is the kind of state for which it makes sense to ask for reasons or justification. This approach would treat both the low-ability subject and the high-ability subject who arrive at a belief by fallacious reasoning as equally responsible and blameworthy. But, I suggested that they are not equally blameworthy. These kinds of considerations may motivate a modification to a simple answerability account and/or an alternative account which can allow that differences in a subject's ability to recognise and respond to reasons makes a difference to the degree of their responsibility and blameworthiness. A range of different accounts of responsibility have been offered on which one is not equally responsible for all states for which one could be asked to provide a

justification or reason. These include McHugh's and McCormick's accounts which make responsibility for belief depend on one's ability to recognise and respond to reasons; Brink and Nelkin's account on which responsibility for X depends on the quality of one's opportunity to do otherwise; and Sher-style accounts of responsibility on which whether one is blameworthy in part depends on whether one fails to meet standards one could have been reasonably expected to meet (Sher 2009, McCormick 2011, Brink and Nelkin 2013, McHugh 2013).¹⁵

In summary, a large range of different approaches in epistemology find it useful to appeal to a notion of blameless/blameworthy belief in addition to traditional notions such as being justified/unjustified or conforming to/violating epistemic norms. Furthermore, the sense in which the relevant beliefs are blameworthy does not appear to be moral. Having defended the utility of the notion of epistemically blameless/blameworthy belief, I go on in the rest of the paper to investigate what is the best account of epistemic blame. In particular, I develop an account of epistemic blame by examining what I take to be the most promising variable account of moral blame—Sher's influential account.¹⁶ The account I will develop of

¹⁵ On Sher's account, S is responsible for doing wrong if S is aware that it is wrong or unaware that it is wrong despite evidence that it is wrong her failure to recognise which is substandard and caused by attitudes, dispositions or traits that are constitutive (143). He holds that what any given agent should recognise in part depends on the agent's capacities (113-15).

¹⁶ Two other influential variable accounts are Scanlonian accounts and functional accounts. According to functionalist accounts, blame is defined as a state which plays a certain functional role (e.g. McGeer 2013). For instance, it may be suggested that to morally blame is to judge the target morally blameworthy and, as a result, to feel and act in ways that have the

epistemic blame bears important similarities to Sher's account yet differs significantly due to the differences between epistemic and moral blame.

4. Sher's account of moral blame.

Sher endorses a "two-tiered" account of moral blame "which takes it to consist of a characteristic set of affective and behavioural dispositions that are organised around a characteristic type of desire-belief pair" (2006:14-15). The relevant dispositions are just those we ordinarily associate with blame, such as the disposition to have certain kinds of negative emotional responses and to behave in certain ways, e.g. to reproach the wrong-doer and apologise, if the wrongdoer is oneself. However, he argues that it would be inadequate to understand blame merely as the relevant dispositions. For, he says this would fail to explain what unifies the various dispositions and so why blame merits a single term or why the corresponding concept is morally important (98). Instead, he suggests that what unifies these various dispositions is that they are the upshot of a certain belief-desire pair, namely the belief that someone has acted badly or has a bad character, and the desire that they not have

function of improving conformity with moral norms. For reasons indicated in the earlier discussion (S3) of instrumentalism about blame, I find such accounts unappealing. According to Scanlon, blaming involves judging that someone has done something that indicates intentions or attitudes that are faulty by the standards of a relationship and, as a consequence, modifying one's understanding of that relationship in a way that the judgement makes appropriate (2008). Like others, I am sceptical that the notion of a relationship can do the work Scanlon needs to explain what it is to be blameworthy and to blame (e.g. Wallace 2011, Bennett 2013, Sher 2013).

acted badly or not have a bad character (103).¹⁷ As he explains, to act badly is not only to violate a moral norm but to do so without excuse or justification (9).

Given the specific content of the belief-desire pair in terms of which Sher provides an account of blame, it is not directly applicable to our cases of blameworthy belief. For belief does not appear to be either an action or character trait. Further, as we saw earlier, in the relevant cases we blame someone for a belief even when they don't seem to have violated a moral norm. However, in the next section I will develop an account of epistemic blame which bears important similarities to Sher's account although differing in important ways in its details. So it's worth exploring Sher's original account and how it deals with objections as a prelude to developing the new account of epistemic blame. For some objections levelled at Sher's account could also be levelled at the later account.

Given that Sher's account treats blame as not merely a belief but also involving a desire and certain dispositions, it can meet the three criteria we earlier set out for an account of blame: it allows that blame goes beyond judging blameworthy; it can be private and it can take various forms. Central to Sher's account is his explanation of how the various dispositions associated with blame are unified by the relevant belief-desire pair. In particular, he suggests that the various dispositions characteristic of blame can be explained by appeal to the known effects of frustrated desire (104-5). In general, frustrated desires dispose one to a variety of reactions. In particular, when a desire is frustrated this disposes one 1) to feel badly about not getting what one wants; 2) to publicly express the unsatisfied desire; and 3) to

¹⁷ While Sher's account allows for blame for both bad actions and bad character, for simplicity I will mostly focus on acting badly in what follows.

substitute the pursuit of some related but more achievable goal.¹⁸ In the case of blame, Sher suggests that the negative emotions typical of blame are an example of how one feels badly when one's desires are frustrated. Further, the expression of blame in reproach or apology can be understood as a public expression of the unsatisfied desire. For instance by apologising for one's past wrongdoing one makes it clear that one now wishes one had not done it. By reproaching some third party for their wrongdoing one makes it explicit that, if it were up to one, they would not have acted badly (109). However, reproach and apology are not merely expressions of unsatisfied desire for one could express one's desire by simple optatives such as, "I wish you hadn't done that" (110). Sher suggests that reproach and apology are further rationalised by the pursuit of a related goal which is more achievable than the frustrated desire that the wrongdoer not have acted badly.

For instance, where one blames oneself for a past bad action, one cannot take the past bad action back. However, a related thing one can do instead is to apologise. This apology expresses the fact that one now appreciates the force of the moral reason which one previously flouted or ignored in acting badly. Similarly, although one cannot now make it the case that another did not do a bad action, a related thing one can do is to reproach them, where reproach aims to bring the wrong-doer to appreciate the force of the moral reason she previously flouted. Thus, Sher suggests that both reproach and apology are ways of achieving the aim that the wrong-doer now appreciate the moral reason they previously ignored or flouted (111). Further, he claims that, if fulfilled, the aim that the wrong-doer now appreciates the relevant moral reason is the best alternative to the frustrated aim that she not

¹⁸ The original desire need not causally produce a second, substitute desire. Instead, the original desire might have been disjunctive or the original desire might directly motivate the substitute behaviour.

have acted badly. Having summarised the elements of Sher's account of moral blame, I first show how it can deal with a number of objections before turning to examine how it can be used to develop an account of epistemic blame.

Smith objects to Sher's account that it mis-classifies certain cases. Consider her example of a loving mother who judges her son blameworthy for a crime and believes that loving him even more is the best way to increase his conformity to moral norms (2013:35). Smith points out that if the mother goes on to shower her son with love she hardly counts as blaming him. But, the mother may have the belief and desire central to Sher's account, namely the belief that the son has acted badly and the desire that he had not done so.

However, in fact Sher's account need not treat the mother as blaming her son since he argues that blame includes not just the relevant belief-desire pair but also its characteristic dispositions. Indeed, he rejects the suggestion that he should identify blame just with the relevant belief-desire pair irrespective of whether it also gives rise to the various characteristic dispositions (112). More broadly, he argues that he is interested in "the nature of blame itself" and what in fact fulfils the causal and normative role we know it occupies (112). He holds that what occupies it is the belief-desire pair which in normal human beings gives rise to the relevant dispositions. From this perspective, he suggests, it is not of central interest what we should say about someone who had the relevant belief and desire but had a psychology that was so non-standard that they lacked any of the affective or behavioural dispositions that normally accompany such desires and beliefs. Blame, as a natural phenomenon in human beings, involves both the relevant belief-desire pair and the characteristic dispositions.

If that's what blame is, then Sher is in a good position to deal with the case of the loving mother understood as lacking the affective and behavioural dispositions characteristic of blame in relation to her son's wrongdoing. Of course, we needn't assume that the mother

lacks the dispositions characteristic of blame in all contexts. For instance, even if the loving mother responds to her son's wrongdoing by showering him with further love, she may feel guilt and reproach herself for her own bad actions. However, if the mother is not disposed to feel angry and reproach her son, then Sher can deny that she blames the son.

Of course, an alternative description of the loving mother is that despite her non-blaming outward behaviour, she in fact experiences the emotions characteristic of blame. For instance, perhaps, the mother does feel indignant about what her son has done and is disposed to show this in her behaviour, but since she thinks that to do so would be counter-productive, she manages to conceal her feelings and instead behaves even more lovingly toward him. In this case, Sher's account would commit him to saying that the mother does blame her son. And, surely, that's the right result. If the mother has the emotional reactions characteristic of blame, but manages to suppress them in her outward behaviour, this looks to be a classic case of private unexpressed blame.¹⁹

¹⁹ A further variant would be one in which the mother is disposed to act and feel in ways characteristic of blame, say being disposed to feel anger and rebuke him, but fails to act and feel that way. On this variant, Sher's account commits him to holding that the mother does blame her son. This might not seem an entirely plausible verdict. However, Sher could reply by saying that it's intuitively unclear what to say about this variant case, and so argue that it's not a major problem that his account treats it as a case of blame. Alternatively, Sher might want to adapt his original account to require that blame involve at least some of the dispositions characteristic of blame being manifest. This would still allow for private blame since the dispositions could be manifested in feelings which are kept private. It would also be compatible with the variability of blame for the suggestion is merely that *some* of the dispositions characteristic of blame be manifest.

Having seen how Sher can deal with the loving mother example, I turn to consider the objection that the relevant desire is not necessary for blame (e.g. Smith 2013). For example, a manager might simultaneously blame an employee, Ed, for some bad act (e.g. sexually harassing a fellow employee, or stealing money from the firm), yet be all too happy that Ed acted badly since it gives them a perfect opportunity to get rid of Ed whom they have wanted to get rid of for years. Analogous cases seem possible in the epistemic case. For instance, a politician might blame his political opponent for believing against the evidence, say, that voters will support tax increases but be simultaneously quite happy since it makes his opponent less likely to win.

Sher might attempt to deal with these putative counterexamples by arguing that, in these cases, the blamer really does have the relevant desire, but it's just not very strong. Perhaps, for example, the employer does have a desire that Ed not have done a bad act, but has a stronger desire to have grounds for firing him. Of course, if the relevant desire is weak, then its non-satisfaction doesn't explain the strength of the employer's negative feelings and actions.²⁰ However, we may reasonably suppose that the employer's other desires explain these. For instance, her desire to fire her employee explains why she does fire them and reproaches them as part of her public justification for doing so.

Even if Sher can deal with the case in which the employer has a weak desire that the employee did not act badly, it may be suggested that there is a possible variant case in which the employer doesn't have even a weak desire that the employee had not acted badly. In

²⁰ Sher says, "no mere wistful yearning could account for the vehemence and persistence of the reactions by which we manifest blame. To explain why we tend to get mad at those we blame, why we often go out of our way to reprimand them, and so on, our blame-constituting desires must be quite strong" (128).

reply, Sher can argue that in cases where the blamer believes that the target acted badly, acts toward the target in ways characteristic of blame but doesn't desire that the target didn't act badly, they are not blaming at all. That the blamer believes the target acted badly doesn't entail that they blame the target since to judge blameworthy is distinct from blaming. That the blamer acts towards the target in ways characteristic of blame is consistent with them not in fact blaming since we should not identify blame with any external action given the possibility of private blame. (Recall that according to our earlier criteria, an account of blame should 1) distinguish judging blameworthy and blaming and 2) allow for private blame.) Furthermore, it is plausible that any account of blame should allow for fake or pretend blame whether in a theatrical setting or even in real life where it is instrumentally useful to make others think one blames when one does not. For instance, a racist might find it useful to hide his racism on occasion by appearing to blame others for their racist behaviour even while he secretly approves of it. To the extent that he can summon up not only characteristic blaming behaviour but also negative emotions, his pretence will be even more believable. So any account should allow that characteristic blaming behaviour even when accompanied by negative emotion toward the target may be mere pretend blame. This is what Sher should say about the employer who judges that an employee has acted badly and acts toward the employee in ways characteristic of blame in order to get rid of them, but lacks the desire that they not have behaved badly: it's just a case of pretend blame.

In conclusion, Sher's account can deal with Smith's loving mother example as well as her objection that desire is not necessary for blame. In the next section, I develop Sher's account to provide an account of epistemic blame.

5. A new account of epistemic blame.

Having defended Sher's account in the moral case, let us now try and extend it to the epistemic. Following his original model, the idea would be that epistemic blame consists in characteristic dispositions to feel and behave which are unified by a characteristic belief-desire pair. In the moral case, the relevant belief is that someone acted badly and the relevant desire is that they had not acted badly.²¹ We cannot simply apply this account to the case of belief since belief is not an example of action. Furthermore, a blameworthy belief need not be bad in the same sense as blameworthy action. For example, when Maud dogmatically believes against the evidence, she doesn't seem to have violated any moral norm or done anything morally wrong. In general, one can violate an epistemic norm such as the norm of conforming one's beliefs to the evidence without violating a moral norm. Relatedly, one can violate a moral norm – say by knowingly poisoning an innocent – without breaking any epistemic norm governing belief. Thus, to provide a Sher-style account of epistemic blame we need to modify the content of the relevant belief-desire pair which figures in his account of blame.

In the case of moral blame for action, the relevant belief-desire pair concern acting badly. By contrast, in the case of epistemic blame, the relevant belief-desire pair concern believing badly. Acting badly goes beyond acting wrongly or violating moral norms. For one can violate a norm while having an excuse or justification. Similarly, in the epistemic case,

²¹ It wouldn't help to appeal to the notion of a bad character trait. For, our judgement that Maud is blameworthy doesn't seem to be a judgement about her character traits. For instance, in judging that Maud is blameworthy, we don't examine whether she is generally disposed to dogmatically ignore the evidence. Indeed, she would be blameworthy even if she is not generally disposed to be dogmatic. So the sense in which she is blameworthy is not that she either has or manifests a bad epistemic character trait.

the target of blame is not merely someone whose beliefs violate epistemic norms. For, one could have an excuse for violating an epistemic norm, say because one had strong but misleading evidence that one was not violating it. On the standard view that epistemic reasons are incommensurable with other reasons, one cannot have a justification for violating an epistemic norm, but only an excuse.²² So the relevant belief-desire pair in the case of epistemic blame would focus on those whose beliefs violate epistemic norms without excuse.

Of course, many agents would not phrase their blame by appealing to such terms as “epistemic norms”. But, we are nonetheless familiar with criticising others and being criticised for violating epistemic norms even if those doing the criticising wouldn’t phrase it by appeal to the expression “epistemic norm”. More likely, they would phrase it in terms of the content of the relevant norm. For instance, many hold that the norm of belief is to conform one’s beliefs to the evidence. We routinely criticise people, ourselves and others, for failing to conform to this norm although we don’t phrase the criticism using the expression “the norm of belief”. Instead, we criticise people for dogmatically believing against the evidence or for failing to believe even in the face of overwhelming evidence. For ease of expression, it will be useful to employ the phrase “believe badly” to cover cases in which a subject’s beliefs violate epistemic norms where the relevant subject has no excuse. So, I will say that, in the case of epistemic blame, the relevant belief is that the agent believed badly

²²Suppose that a demon credibly promises that if you believe that p against the evidence you will obtain great riches or, alternatively, prevent a moral catastrophe. If epistemic reasons are incommensurable with others then one has no all things considered reason to believe that p. Rather one has an epistemic reason not to believe that p and an incommensurable prudential or moral reason to believe that p.

and the relevant desire is that the agent hadn't believed badly. This allows us to remain neutral on the philosophically controversial issue of what the norm of belief requires.

Just as in Sher's account of moral blame, the relevant belief-desire pair unifies the affective and behavioural dispositions characteristic of epistemic blame. Just as in the moral case, epistemic blame involves a frustrated desire. However, in the epistemic case, the relevant desire has a different content: that the agent had not believed badly. Just as in the moral case, frustrated desire leads to negative feeling and related behaviour. For instance, when we blame Maud for dogmatically believing against the evidence that the President is in New York, the frustrated desire that she not believe badly leads to negative feelings such as our anger with her. Furthermore, this frustrated desire is publicly expressed in rebuke when we say such things as, "How can you believe the president is in New York given that the TV is showing that he is in Washington?!" As in the moral case, the frustrated desire also leads to the pursuit of an alternative more easily satisfied desire. Whereas in the moral case, the alternative desire was the desire that the wrongdoer appreciate the moral reasons she previously ignored or flouted, in the epistemic case the substitute desire is that the believer appreciate the relevant epistemic reason that she previously ignored or flouted. This alternative desire also explains why we rebuke Maud in an attempt to get her to appreciate the evidence she previously ignored or flouted.

While Sher's account of moral blame can be extended to provide a plausible account of epistemic blame, it may be thought to face the objection that people don't have the relevant desire which is partly constitutive of epistemic blame, namely the desire that the believer not have believed badly. It's standard in the epistemic literature to raise doubts about whether believers generally desire that their beliefs meet the norm of belief. Plausibly, if we don't desire that believers conform to the norm of belief then we also don't desire that believers don't violate the norm of belief without excuse. But if many of us don't have the

desire that others don't believe badly, then the Sher-inspired account of epistemic blame faces the objection that many of us cannot epistemically blame others for we lack the relevant desire.

For instance, on the supposition that the norm of belief is either truth or knowledge, it's been questioned whether believers have a desire for true belief or knowledge about any matter whatsoever. For instance, it's not clear that we desire to believe all trivial truths such as the truth about the number of grains of sand on the local beach. Indeed, we are constantly passing up opportunities to gain true beliefs at very little cost, say by browsing articles on Wikipedia (Kelly 2003). Thus, many question whether we have a general desire for true belief or knowledge about any matter whatsoever. Further, certain examples suggest that agents sometimes have a desire not to have a true belief or knowledge about some matter. For example, consider a birthday girl who wants a surprise birthday present and so doesn't want to have a true belief about what her present is. Indeed, even if she finds herself with a belief about what her present is as a result of happening across some evidence, she may desire that her belief is false so that she still has a surprise on the day. These kinds of considerations have led many epistemologists to doubt that we have a general desire for true belief/knowledge about any matter whatsoever, or at least to doubt that we have a strong desire for true belief/knowledge about any matter whatsoever.

While the epistemic literature has been concerned with whether believers have a desire that they themselves have true beliefs/knowledge, the desire relevant to the proposed account of epistemic blame is not merely a first-personal desire. For, of course, we blame not just ourselves but also others for violating epistemic norms with no excuse. So, the desire required by the proposed account of epistemic blame is that both oneself and others don't believe badly. But, it's even easier to come up with reason to doubt that we have the relevant third-person desire than the first-person desire. We constantly pass up opportunities to give

others true beliefs or correct their false beliefs, even when we can do so at little cost. For instance, in the course of a conversation with a colleague I may realise that she fails to believe a number of truths about matters of little significance, such as about my cat's name, or where I was born and so on. Indeed, I may realise that she has false beliefs about these matters. While I could provide her with true beliefs at very little cost, I most likely won't bother to do so unless her failure to have the relevant true beliefs will foreseeably lead her into difficulty, or they are the topic of the conversation. Furthermore, given that one's own interests may diverge from those of others, it's routinely the case that an agent may desire that others not have true beliefs or knowledge about some matter. So, for example, during negotiations to sell a house, the seller may desire that the buyer not have true beliefs about the minimum price she will accept for the house, and the buyer may desire that the seller not have true beliefs about the maximum she is prepared to pay.

It seems, then, that there are good reasons to doubt that people have strong desires that they and others have true beliefs/knowledge about every proposition whatsoever. However, this need not cast doubt on the proposed account of epistemic blame. The account doesn't require that each of us, for all propositions and times, has a strong desire that ourselves and others have true beliefs. The account does have the consequence that it's a necessary condition of S1's blaming S2 for believing badly in believing that p that S1 then desires that S2 didn't believe badly in believing that p. But this does not require that S1 then have the general desire that no subject believe badly concerning any proposition whatsoever. Nor is it required that S1 have a strong desire that S2 not believe badly in believing that p when she blames S2 for believing that p. A strong desire is required only in cases in which S1's blame of S2 explains strong emotions and rebuke. For, the frustration of a weak desire would not plausibly lead to strong negative emotions or rebuke (see footnote 20).

Indeed, Sher accepts that the strength of the desire in the moral case varies across persons and times. On his account, the variation in the strength of the relevant desire explains the presence and strength of blame. Sometimes even when subjects are blameworthy, we don't blame them because at that time we lack the relevant desire, or it is present only in a weak form. For instance, a mafia boss may fail to blame his subordinates who are morally blameworthy for their extortion and torture racket because he lacks the desire that they did not act badly. At other times, we blame wrongdoers, but in rather mild ways which are insufficient to the extent of their wrongdoing. Sher would explain this by saying that the blamer has the relevant desire, but it's not very strong. For instance, a troop commander may blame his subordinates mildly for torturing civilians when this provides useful information for the war effort. Thus, it's no problem for Sher's account of moral blame that it's not the case that all subjects at all times strongly desire that they and others don't act badly.

Analogously, it's no problem for the suggested account of epistemic blame that it's not the case that all subjects at all times strongly desire that they themselves and others don't believe badly. Indeed, it's a happy result for the proposed account of epistemic blame if a strong desire that people don't believe badly is less widespread than the strong desire that people don't act badly. For, it seems part of the phenomena to be explained that blaming people for their beliefs is less common and often takes a less strong form than blaming people for their bad actions. This is explicable if the desire that oneself and others don't believe badly is less widely and strongly held than the desire that oneself and others don't behave badly.

6. Differences revisited.

Having shown how Sher's account can be extended to provide an account of epistemic blame, it is useful to revisit some of the differences between our reaction to epistemic and moral failings. As we saw earlier, moral blame is often associated with certain negative

emotions such as resentment; and certain actions such as punishing, and seeking apology and/or compensation. In the case of self-blame, moral blame is associated with feeling guilty, and seeking to apologise or recompense. By contrast, epistemic failings aren't associated with resentment, the demand for compensation or punishment. Nor would it seem appropriate for the subject of the failing to feel guilty, apologise to others or recompense them.

Within the broad approach to blame which I am employing in this paper which allows for the variability of blame, these differences don't show that there can be no genuine epistemic blame. For variable accounts of blame allow that blame needn't involve punishment, demands for compensation, or resentment. However, it's a benefit of the proposed account of epistemic blame than it can explain some of the similarities, as well as the differences, between epistemic and moral blame. On the proposed account, both epistemic and moral blame involve the belief that the target has done badly, where that involves violating a norm without excuse or justification. Furthermore, both involve the desire that the target had not done badly. In each case, the frustration of the latter desire generates characteristic affective and behavioural dispositions. In both cases, these characteristic dispositions include the disposition to be angry and to rebuke the target. However, the dispositions are not identical. In particular, some of the dispositions characteristic of moral blame are not characteristic of epistemic blame, such as the dispositions to punish the target, feel resentment, and demand compensation or apology. Thus, the characteristic affective and behavioural dispositions of moral and epistemic blame are overlapping rather than identical.

The similarities and differences in the characteristic dispositions associated with moral and epistemic blame can be explained by appeal to the content of the relevant belief and desire and, in particular, how the notion of doing badly is to be understood. In the moral case, blame involves the belief that the subject has acted morally badly and the desire that she

did not act morally badly. In many, if not all cases, acting morally badly involves wronging someone. For example, if you intentionally harm me without justification or excuse, you have wronged me. It's for this reason that I resent you, feel hostile to you, and demand compensation and apology. By contrast, when Maud dogmatically believes against the evidence that the President is in New York, it's not clear that she's wronged anyone. As a result, when we blame her, we may be angry and rebuke her, but would not resent her and it doesn't make any sense to demand compensation or that she apologise to someone wronged.

Examining the content of the substitute desire reinforces these differences. In the moral case, the original frustrated desire that the agent didn't act badly leads to the more achievable substitute desire that the target appreciates the reason she flouted or ignored in acting badly. In the moral case, this would involve appreciating a flouted moral reason, say that another's interests deserve to be taken into consideration. The wrongdoer can demonstrate that she now appreciates this flouted moral reason by apologising, or perhaps even compensating the person she wronged. By contrast, in the epistemic case, the flouted reason is epistemic and does not involve ignoring or violating another's interests. As a result, apology or compensation would seem inappropriate in demonstrating that she now appreciates this flouted epistemic reason. Instead, it seems enough to acknowledge the original fault as well as to take steps to reduce the possibility of its future recurrence.

Of course, there can be mixed cases in which a blameworthy belief is both epistemically and morally blameworthy. For instance, a developer who believes against the evidence that the tower block cladding is safe leading to the death of residents by fire is both epistemically and morally blameworthy in believing as he does. Furthermore, he arguably does wrong the residents in believing as he does, for he fails to take proper account of their interests. Thus, in such mixed cases, we would expect the residents would feel anger and hostility towards the developer and demand compensation and apology.

I've now defended an account of epistemic blame modelled on Sher's account of moral blame. In each case, I suggested that blame consists in a belief-desire pair together with the related characteristic set of affective and behavioural dispositions. However, as we have seen the content of the belief-desire pair differs and this explains why the set of affective and behavioural dispositions characteristic of epistemic blame overlaps with, but is not identical to, the set characteristic of moral blame. One way to understand the proposed account of epistemic blame would be to see it as providing an account of a kind of blame distinct from moral blame but which nonetheless shares its general structure as composed of a belief-desire pair together with accompanying characteristic dispositions. However, an alternative would be to try and understand each of epistemic and moral blame as instances of a more general notion of blame which can be triggered by different considerations on different occasions. To flesh out this second alternative one would need to provide a detailed account of this general notion of blame operative in both the moral and epistemic cases. For instance, it might be proposed that both moral and epistemic blame involve a single belief-desire pair together with accompanying characteristic dispositions. For example, since both moral and epistemic blame seem to involve a norm violation without justification or excuse, it might be suggested that the relevant belief-desire pair are the belief that a subject violated a norm without justification and excuse together with the desire that the subject had not done so. The idea would be that this single notion of blame could be triggered by different considerations on different occasions, sometimes by the violation of a moral norm and sometimes by the violation of an epistemic norm. As a result, the proposed general notion of blame could properly acknowledge that different norms are violated in different examples of blame. In particular, it could acknowledge that someone can violate a moral norm in a blameworthy way without violating an epistemic norm, and vice versa. For instance, someone might violate a moral norm in a blameworthy way by knowingly poisoning an

innocent without violating any epistemic norm regarding forming or acting on one's beliefs. Likewise, someone might violate an epistemic norm in a blameworthy way without violating a moral norm as in our case of Maud who dogmatically believes against the evidence that the President is in New York.

For my purposes here, we need not settle the issue of whether there is a univocal account of blame covering both epistemic and moral blame, or whether it is better to regard epistemic and moral blame as different but related kinds of blame. Doing so would involve a much larger discussion of all the different kinds of norm violation there are, the extent and strength of our desire that they are not violated, and whether the frustration of our desire that they are not violated generates affective and behavioural dispositions of such a nature that it is appropriate to classify our response to frustrations of the relevant desire as blame. For my purposes, I need not settle these issues here. My ambition has just been to develop and defend an account of epistemic blame. To do so I need not settle the further issue about whether moral and epistemic blame are different kinds of blame or whether they can ultimately be accommodated by appeal to a more general notion of blame triggered by different considerations on different occasions.

7. Conclusion.

Many hold that we can be blameworthy for our beliefs. Indeed, that we can be so blameworthy is a central theme in much recent epistemological work. If we can be blameworthy for our beliefs, it can be fitting to blame subjects for their beliefs. However, I have argued that the most promising accounts of moral blame cannot be straightforwardly applied to epistemic blame since the latter need not involve a violation of moral norms. Given that, I have provided a new account of epistemic blame inspired by Sher's account of moral blame. The proposed account of epistemic blame has the same overall structure as Sher's

account of moral blame: it involves characteristic affective and behavioural dispositions unified by a characteristic belief-desire pair. However, in the case of epistemic blame the relevant belief-desire pair have a different content. I've argued that this difference in content explains why epistemic blame differs from moral blame in its typical strength, and in the nature of its characteristic affective and behavioural dispositions. With this new account, we can defend the frequent appeal to the notion of blameworthy and blameless belief in contemporary epistemic discussion, including in debates concerning higher-order evidence, peer disagreement, closure, externalism versus internalism, and the norms of belief.²³

Baker-Hyatt and Benton. 2015. "Defeatism defeated". *Philosophical Perspectives* 29, 1:40-66.

Beddor. 2015. "Process reliabilism's troubles with defeat". *Philosophical Quarterly* 65, 259: 145-59.

Bennett. 2013. "The expressive function of blame". In Coates and Tognazzini 2013: 66-83.

Brink and Nelkin. 2013. "Fairness and the architecture of responsibility". *Oxford Studies in Agency and Responsibility* 1:284-313.

Clarke, McKenna and Smith (eds.). 2015. *The nature of moral responsibility*. OUP: Oxford.

Coates and Tognazzini (eds.). 2013. *Blame*. OUP: Oxford.

²³ Thanks to the Mind Association for a Mind Senior Research Fellowship during which I completed this paper. Earlier versions of this paper were presented at the St Andrews Blame Workshop, The Joint Session of the Mind and Aristotelian Society (Oxford 2018) and the European Epistemology Network meeting (Amsterdam 2018). Thanks to audiences on those occasions for helpful comments, and thanks for feedback on the written paper from Matthew McGrath and Brian Weatherson.

- Darwall. 2006. *The second-person standpoint*. Harvard University Press.
- Goldberg. 2015. *Assertion: on the philosophical significance of assertoric speech*. OUP: Oxford.
- Goldberg. Forthcoming. "Should have known". *Synthese*.
- Greco. 2010. *Knowledge as an Achievement*. Cambridge: Cambridge University Press.
- Hawthorne and Stanley. 2008. "Knowledge and action". *Journal of Philosophy* 105:571-90.
- Hawthorne and Srinivasan. 2013. "Disagreement without transparency: some bleak thoughts." In Christensen and Lackey (eds.), *The epistemology of disagreement*, OUP: Oxford, 9-30.
- Hieronymi. 2008. "Responsibility for believing". *Synthese* 161, 3:357-73.
- Hieronymi. 2014. "Reflection and responsibility". *Philosophy and Public Affairs* 42,1: 3-41.
- Kelly. 2003. "Epistemic rationality is instrumental rationality: a critique". *Philosophy and Phenomenological Research* LXVI, 3:612-640.
- Lasonen-Aarnio. 2010. "Unreasonable knowledge". *Philosophical Perspectives* 24.
- Lasonen-Aarnio. 2014. "Higher-order evidence and the limits of defeat". *Philosophy and Phenomenological Research* LXXXVIII, 2: 314-345.
- Lasonen-Aarnio. Forthcoming. "Enkrasia or evidentialism". *Philosophical Studies*.
- Levy. 2005. "The good, the bad, and the blameworthy". *Journal of Ethics and Social Policy* 1, 2:2-16.
- Littlejohn. Forthcoming. "A plea for epistemic excuses". In Dorsch and Dutant (eds.), *The New Evil Demon*, Oxford: Oxford University Press.
- McCormick. 2011. "Taking control of belief". *Philosophical Explorations* 14, 2:169-183.
- McHugh. 2013. "Epistemic responsibility and doxastic agency." *Philosophical Issues* 23.
- McGeer. 2013. "Civilising blame". In Coates and Tognazzini 2013, 162-188.

- Nottlemann. 2007. *Blameworthy belief: a study in epistemic deontology*. Springer: Dordrecht.
- Owens. 2000. *Reason without Freedom*. Routledge.
- Peels. 2017. *Responsible belief*. OUP: Oxford.
- Scanlon. 1998. *What we owe to each other*. Harvard University Press: Cambridge, Mass.
- Scanlon. 2008. *Moral dimensions: permissibility, meaning, blame*. Harvard University Press: Cambridge Massachusetts.
- Sher. 2006. *In praise of blame*. OUP: Oxford.
- Sher. 2009. *Who knew? Responsibility without awareness*. OUP: Oxford.
- Smith. 2005. "Responsibility for attitudes: activity and passivity in mental life". *Ethics* 115:236-271.
- Smith. 2013. "Moral blame and moral protest". Coates and Tognazzini 2013: 27-48.
- Smith. 2015. "Responsibility as answerability". *Inquiry* 58, 2:99-126.
- Sosa. 2014. *Judgement and agency*. Oxford University Press: Oxford.
- Srinivasan. 2015. Normativity without Cartesian privilege. *Philosophical Issues* 25:273-299.
- Sutton. 2007. *Without Justification*. MIT press: Cambridge, Massachusetts.
- Wallace. 1994. *Responsibility and the Moral Sentiments*. Harvard University Press: Cambridge, Massachusetts.
- Wallace. 2011. "Dispassionate opprobrium: on blame and the reactive sentiments". In Wallace, Kumar and Freeman (eds.), *Reasons and recognition: essays on the philosophy of TM Scanlon*, Oxford University Press: Oxford
- Watson. 1996. "Two faces of responsibility". *Philosophical Topics* 24, 2:227-248.
- Weatherson. Forthcoming. *Normative Externalism*.
- Williamson. 2000. *Knowledge and Its Limits*. OUP: Oxford.

Williamson. 2013. "Response to Cohen, Comesaña, Goodman, Nagel and Weatherson on Gettier cases in epistemic logic". *Inquiry* 56:77-96.

Williamson. 2014. "Very improbable knowing". *Erkenntnis* 79:971-999.

Williamson. Forthcoming. "Justifications, excuses, and sceptical scenarios". In Dorsch and Dutant forthcoming.
