

There is a Distinctively Epistemic Kind of Blame

Abstract: Is there a distinctively epistemic kind of blame? It has become commonplace for epistemologists to talk about epistemic blame, and to rely on this notion for theoretical purposes. But not everyone is convinced. Some of the most compelling reasons for skepticism about epistemic blame focus on disanalogies, or asymmetries, between the moral and epistemic domains. In this paper, I defend the idea that there is a distinctively epistemic kind of blame. I do so primarily by developing an account of the nature of epistemic blame. My account draws on a prominent line of theorizing in moral philosophy that ties blame to our *relationships* with one another. I argue that with my account of epistemic blame on hand, the most compelling worries about epistemic blame can be deflated. There is a distinctively epistemic kind of blame.

1. Introduction

Moral blame is a central part of our private and public lives. Because it is so central, a great deal of philosophical work has been done on the nature and norms of moral blame. Relatively recently, epistemologists have begun focusing more attention on something they call ‘epistemic blame’ (Brown forthcoming, 2017; Rettler 2017; Nottelmann 2007; Peels 2016; Williamson forthcoming; Littlejohn forthcoming).¹ Epistemic blame is supposedly distinct from moral blame, while nevertheless sharing some of moral blame’s characteristic features. For example, it is a response made fitting by a normative failing, a response that goes beyond mere negative evaluation or grading. Epistemic blame has a characteristic ‘sting, force, and depth’ (Hieronymi 2004). But it is not a *moral* response. Epistemic blame has something specifically to do with epistemic failings.

Is there a distinctively epistemic kind of blame? Not everyone is convinced. Some of the most compelling reasons for skepticism about epistemic blame focus on disanalogies, or asymmetries, between the moral and epistemic domains. One potentially salient difference, for example, concerns what I call the comparative ‘coolness’ of epistemic judgment. Consider judging that Bob is epistemically unjustified in

¹ Some historical writers also seem sympathetic to epistemic blame, or something in the neighborhood: ‘[I]f I affirm what is not true, it is evident that I deceive myself; even though I judge according to truth, this comes about only by chance, and I do not escape the blame of misusing my freedom...’ (Descartes 1955:176). Locke and Clifford are also often referred to as important figures in the history of ‘epistemic deontology’; it is not difficult to find passages in their work where they appear to be countenancing epistemic blame (Owens 2000; Plantinga 1993).

believing that there are 799 motes of dust on his desk, or that Tim ought to know how many states there are in the U.S. It seems one can make such judgments, and a wide range of similar others, without feeling any sort of emotion or concern for the matter. Meanwhile, it is perhaps more difficult to imagine making a moral judgment in a similarly cool fashion. Of course, there may be comparably trivial cases in the moral domain. And not all epistemic judgments concern trivial matters. But the coolness of epistemic judgment seems to be a characteristic feature of many forms of epistemic judgment. Perhaps this observation is evidence that talk of distinctively epistemic blame is somehow misguided. If part of what makes the sting, force, and depth of moral blame appropriate is the heatedness of the normative judgments that go along with it (or partly comprise it), and if epistemic judgment characteristically lacks that feature, then there may be something mistaken about the idea of epistemic blame.

This line of thought is suggestive, but we need more details. Whether or not it's plausible depends, for one thing, on our understanding of the nature of blame. Perhaps blame is compatible with cool normative judgment. Who knows? At the very least, we need a better understanding of what blame *is* before we can answer this question. In this paper, I set out to defend the idea that there is a distinctively epistemic kind of blame. I do so, primarily, by developing an account of the nature of epistemic blame. My account draws on a prominent line of theorizing in moral philosophy that ties blame to our *relationships* with one another (Strawson 1962; Scanlon 1998, 2008). I extend this idea, and argue that epistemic blame is a kind of *relationship modification*. Making this extension requires giving a detailed account of what an 'epistemic relationship' consists in. In slogan form, my view is that an epistemic relationship consists in intentions, expectations, and attitudes people have towards one another which are oriented towards their epistemic agency in distinctive ways. I develop an account of what I call the 'general epistemic relationship'. By thinking of epistemic blame in terms of a kind of modification to this relationship, we can do justice to the thought that a characteristic feature of epistemic evaluation is that it is indeed often cooler than moral evaluation, while nevertheless making room for an important distinction between merely negative epistemic evaluation, and epistemic blame.

In Section 2, I formulate the skeptical idea above in more detail. I call it 'the problem of cool judgment'. In section 3, I argue that a relationship-based approach to the nature of blame has resources that will be helpful for developing a solution to the problem. In section 4, I show how we can extend the relationship-based approach to the epistemic domain, presenting my account of the general epistemic relationship. In section 5, I defend relationship-based approaches to blame against a number of objections. In Section 6, I consider a few other reasons we might be skeptical of epistemic blame. I argue that, with

the relationship-based account of epistemic blame on hand, none of them has much force. There is a distinctively epistemic kind of blame.²

2. The Problem of Cool Judgment

Recently, R.J. Wallace (2013) has argued that the sting, force, and depth of blame are to be explained in terms of the idea that blame is essentially a kind of *negative emotional response*. Blame essentially involves negative reactive attitudes such as indignation, resentment, and guilt (in the case of self-blame). Wallace is not alone in this view. Indeed, he is a prominent voice among a large number of moral philosophers who take a broadly ‘Strawsonian’ approach to blame (Strawson 1962; Wolf 2013; Menges 2017). Why take such an approach?

Consider the familiar distinction between judging that something is *valuable*, and actually *valuing* it. Wallace proposes we model the distinction between judging blameworthy and blaming on this distinction. The difference between judging X valuable and actually valuing X comes down to a difference in one’s *engagement* with X. When one values X, one is engaged and exercised by X. One gets excited by opportunities to X, disappointed by missing opportunities to X. Meanwhile, in judging X valuable, one merely recognizes that there are reasons for someone to be so exercised, without yet necessarily being that way oneself.

According to Wallace, something similar is true of blame. When one blames someone for doing something, one is engaged and exercised by what they have done. It matters to one that this person has done something bad, or flouted some norm, and by blaming them for it one manifests this concern in an engaged way. Meanwhile, in judging someone blameworthy, one merely judges that there are reasons for one (or someone else) to be so engaged, without yet necessarily being that way oneself.

The key point, according to Wallace, is that the *emotions* are an integral part of what it is for human beings to be ‘engaged’ and ‘exercised’ in this way. Imagine someone claiming to value opera without feeling any emotion whatsoever about the prospect of going to the Met. Other things being equal, it would be natural to conclude that they do not actually value opera (even though they might judge it valuable). Similarly, according to Wallace, someone who claims to blame someone for a wrongdoing without feeling any emotion would raise suspicion: it would be natural to judge that they are not actually engaged by the norms they claim to be engaged by, and so do not actually *blame* the person, though they may nevertheless judge them blameworthy.

² I should acknowledge that this paper simply assumes a stance in the debate on the ethics of blame that blaming others (and oneself) can at least sometimes be a justified response to a wrongdoing. Some have defended the view that blame is never justified, or that we should seek to modify our blaming practice, for example, in favour of other forms of holding responsible that do not involve blame. See Gokhale (2019) for helpful discussion.

This is not yet an argument that the specific negative emotions of *resentment*, *indignation*, and *guilt* are essential to blame. To arrive at the claim that resentment, indignation, and guilt are the uniquely fitting responses to wrongdoing, the ones that show a kind of engagement with morality in the relevant way, Wallace appeals to the ‘relational character’ of moral norms (Wallace 2013: 369). The basic idea is that blame is a response to violations of norms that constitute (at least part of) our relationships with one another. Caring about these relationships means caring about being treated with mutual regard and respect. According to Wallace, the distinctively reactive sentiments are the fitting response to failures to be treated that way. Thus, Wallace concludes:

[...T]he tendency to blame can be seen to be a peculiarly appropriate way of taking to heart the values around which morality is structured. It matters to us whether or not people succeed in relating to us (and to each other) on moral terms, and this shows itself in the reactive emotions that I have contended are essential to moral blame (369).

How might all of this figure in an argument against epistemic blame?

While some epistemic judgments may have significance for us, and be far from completely cool in nature, there is nevertheless something plausible to the claim that negative emotional responses such as indignation and resentment are out of place in the epistemic domain. Insofar as we keep the distinctively epistemic dimension of someone’s failure in mind, negative emotional responses such as indignation and resentment do not really seem like typically fitting responses. The two examples from the introduction help illustrate this idea (motes of dust, and States in the US). Consider also how you might respond to a colleague upon learning that they’d formed a belief about Prince Harry’s love life, merely on the basis of the tabloids they read; or think of how you might respond to a friend for believing that they’ve got a winning scratch-n-win lotto ticket, just because they have not won in quite some time. While the details of the cases require filling out, it’s hard to imagine what would make indignation and resentment fitting responses *to these epistemic failings*.³ And yet, these sorts of cases are typical examples of what epistemologists have in mind when they talk about epistemic blame. If Wallace is correct, and blame is essentially linked to emotion in the way he has argued, then the idea of distinctively epistemic blame seems misguided. ‘Blame’ would somehow be just too strong of a term for the ways in which we go about criticizing or otherwise negatively evaluating each other for typical epistemic failings.

Of course, there are other approaches to the nature of blame, and not all of them put negative emotional responses at the center of their accounts. According to Angela Smith, for example, what is

³ Perhaps indignation and resentment *can* be fitting responses to such failings. If so, that seems like good news for my main thesis, which is that there is a distinctively epistemic kind of blame.

distinctive of blame is that it is a response that necessarily involves the disposition to register *protest* in light of one's judgment that one has been *wronged* by another (2013: 36). Pamela Hieronymi argues that blame necessarily involves the judgment that the target of blame has shown the blamer (or someone else) *ill will*, and reacting accordingly (2004: 116). In another prominent approach, George Sher argues, roughly, that blame is the manifestation of a frustrated desire that the target of blame not have acted badly (2006: 101-108).

Each of these approaches has its own explanation of what it is about blame that goes beyond mere negative evaluation. What is important for present purposes is that, while none of them gives as central a role to the negative moral emotions as Wallace, it's not clear that the features they appeal to in explaining the significance of blame would be any more at home in the epistemic domain than negative emotional responses. Registering protest in light of a judgment that one has been *wronged* by another is not a fitting response towards that person for fallaciously reasoning about their scratch n' win lotto ticket. The same point seems true with respect to judgments that the target of blame has shown one *ill will*. It seems odd to think of many of the most typical kinds of epistemic failings frequently presented as epistemically blameworthy failings (fallacious reasoning, wishful thinking, intellectual laziness) as involving one person wronging another, or showing another person ill will.⁴ Perhaps Sher's account is more applicable in this regard.⁵ But if indignation and resentment are out of place in the epistemic domain, I would argue that so are judgments of ill will and protestations of wrongdoings.

The problem of cool judgment is not merely an upshot of Wallace's view about the nature of blame. There appears to be a more general tension between the observation that blame is a kind of response that goes beyond mere negative evaluation—that it has a certain sting, force, and depth—and the observation that epistemic judgment is a typically cool (or cooler) sort of normative judgment than moral judgment. Does this tension suggest that there is no such thing as a distinctively epistemic kind of blame?

In the next few sections, I argue that it does not. I focus on an approach to the nature of blame that provides resources for accounting for the significance of epistemic blame, without running roughshod over the observation that epistemic judgment is a comparatively cool form of normative judgment.

⁴ Below I discuss perpetrators of epistemic injustice as possible targets of epistemic blame. Perhaps such cases are ones in which a judgment that someone has shown you ill will, or at least wronged you, would be fitting. See discussion of epistemic injustice in Sec. 4.

⁵ Jessica Brown has argued that Sher's account extends naturally to the epistemic domain (Brown forthcoming; see Author MS for a response to Brown's argument). In my view, even the manifestation of *frustrated desires* that people not have acted (or, believed) badly seem somewhat out of place in the epistemic domain.

3. Blame and Relationships

A number of prominent moral philosophers would agree with Wallace that blame has something important to do with relationships. However, some have argued that we can appeal more directly, and more broadly, to certain kinds of *relationship modification*, in order to account for the significance of blame (Scanlon 1998; 2008). Since Scanlon is perhaps the most prominent of these voices, I focus on some of his core ideas.

Scanlon works with a somewhat technical definition of ‘relationship’. For Scanlon, our relationships are constituted by certain intentions, expectations, and attitudes we have concerning how to act and feel towards one another. Strictly speaking, relationships also consist in actual dispositions people have towards each other (in addition to their *intentions*, for example), as well as the *reasons* why they have those intentions and dispositions. I will bracket these additional points for the sake of presenting the basic idea of the account.

Distinguish between relationships as normative ideals—idealized sets of intentions, expectations, and attitudes that would be constitutive of a *good* relationship, whether friendships, romantic relationships, or something else—and token relationships; the latter are specific instances of actual relationships which will approximate to a better or worse degree the relevant normative ideal. Consider the normative ideal of friendship. Paradigm intentions and expectations of friendship might include ‘intending to give help and support when needed, beyond what one would be obligated to do for just anyone; intending to confide in the person and to keep his or her confidences in return; and intending to spend time with the person when one can, and to “keep in touch”’ (Scanlon 2008: 132). They might include the expectations that your friend will do the same. In addition to a central core of intentions and expectations that one must have in order to count as a good friend, there are plausibly also non-obligatory intentions that are nevertheless good to have according to the normative ideal. Perhaps intending to take pleasure in the successes of one's friend, or to generally hope that things go well for them, are two such examples.

A judgment of blameworthiness is a judgment (perhaps implicit) that someone with whom you stand in some relationship has intentions, expectations, or attitudes that in some way fall short of the normative ideal of that relationship. My friend Terrence promised to pick me up at the airport on Tuesday, and he never showed up. Because he didn't have a good excuse, I judge Terrence blameworthy—I judge that Terrence has certain intentions, expectations, or attitudes towards me that in some way fall short of the normative ideal of friendship. His failing to show up after promising to do so says something about his regard for me as a friend. Friends keep their promises.

A *blame* response would then consist in a modification to the intentions, expectations, and attitudes I have towards Terrence, in a way made fitting by the judgment that he is blameworthy. So, if I

blame Terrence for not picking me up, according to Scanlon, this might simply be a matter of modifying my intention to take Terrence at his word in the future, and/or modifying my expectations that he will keep his word in the future; depending on the severity of the breach of confidence, it might be fitting to stop intending to feel pleasure at Terrence's successes in the future, for example, or to cut off relations altogether. This combination of judgments, and modifications to intentions and expectations as a result, is what my blame consists in.⁶

Notice that reactive attitudes such as indignation and resentment are allowed a role in our understanding of blame on this account. However, they are not central. Forming certain reactive attitudes under certain circumstances may constitute *one* way of modifying our relationships with one another. But relationship modification can be done in many ways. Indeed, as I've noted, Scanlon's view belongs to a broad family of views including Peter Strawson's influential account of blame (1962). Both Scanlon and Strawson understand blame in terms of relationship modification—however, they disagree about what kinds of modification are relevant. For Strawson, like R.J. Wallace, only a fairly restrictive range of attitudes including indignation and resentment are relevant. For Scanlon, it is much broader, with an emphasis on changes in intentions and expectations.

The relationship-based approach to blame is not without critics. A prominent objection—due in independent ways to R.J. Wallace (2011) and Susan Wolf (2011)—is that the relationship-based approach 'leaves the blame out of blame' (Wallace 2011: 349). While proponents have advertised as a *feature* of the approach that it can accommodate the coolness of certain kinds of blame, Wallace and Wolf argue that it goes too far in this regard. According to Wallace, as we have seen, reactive attitudes such as indignation and resentment simply are essential to what it really means to blame someone. The fact that Scanlon's account does not give them a central role is a defect of the account, according to Wallace (and Wolf).

Ultimately, I do not want to take a stand on whether this is true of moral blame.⁷ But I will emphasize two points. First, for those who are sympathetic to negative emotional response accounts of blame, it is important to remember that we can allow that certain kinds of emotion will normally accompany blame-responses. Second, this objection reveals one reason why an extension of the relationship-based account to the *epistemic* domain seems particularly promising. As I have argued, while

⁶ An anonymous referee has pointed out that the account may strike some as implausibly *voluntaristic* about blame. In particular, it seems an important feature of our blaming practices that blame is a response that often *happens* to us, in the way that we may find ourselves feeling anger or resentment towards another person, rather than deciding to feel anger or resentment. In response, note that the account does not state that modifications to one's intentions, expectations, or attitudes must be something over which the blamer has direct voluntary control. Indeed, on my interpretation of the account, we can leave it open whether this sort of thing may happen to one, in the same way that one may find oneself feeling anger or resentment towards another person. Thanks to an anonymous referee for pressing this issue.

⁷ It does seem to me that by maintaining that negative emotional responses are essential to blame, we run into a powerful worry about the possibility of affectless blame (Sher 2006, 88; Fricker 2016, 170)

reactive attitudes such as indignation and resentment may be integral to moral blame, they do not seem integral to epistemic blame. The account's broader, more direct approach to relationship modification is precisely what underpins its promise for avoiding the problem of cool judgment. If the approach can be extended to the epistemic domain, we may be in a position to understand the significance of epistemic blame in terms of a kind of relationship modification that does not seem out of place in the epistemic domain.

Other prominent objections (e.g. Smith 2013) focus on ways in which the account seems forced to misclassify cases. I will return to some of these in Section 5. For now, I mention just one such case, since it will help clarify some terms at the heart of the approach. Assume for the sake of argument that the normative ideal of a parent-child relationship involves children caring for elderly parents in certain ways. Assume that, knowing that no-one is perfect, Rosemary arranges a backup plan for her old age to ensure she has care even if her son fails to *fully* live up to the ideal in question. It seems she judges that her son falls short of the normative ideal of the parent-child relationship, and she modifies her expectations and intentions in a way made fitting by this judgment. Yet, it doesn't look like she thereby blames her son.⁸

In response, note that the notion of a relationship as 'normative ideal' is not meant to imply perfection on behalf of any party to that relationship. Rather, the normative ideal of the parent-child relationship specifies intentions, expectations, and attitudes characteristic of a *good* parent-child relationship—it specifies intentions, expectations, and attitudes that parties to the relationship ought in this sense to have. We can spell this out in such a way as to account for the fact that no-one is perfect. It seems to follow from this that it would be inappropriate to judge that your child falls short of the normative ideal of the parent-child relationship, simply because you judge that they may not be perfectly reliable in taking care of you in your old age. For this reason, the relationship-based account does not misclassify the case.

A more general note before moving on. For those who are already convinced that relationship-based approaches to moral blame are mistaken, I encourage reading this paper as a whole in the following way. Rather than reading the paper as presupposing the truth of relationship-based approaches to blame (and then applying them to the epistemic domain), it can be read as suggesting that the fruitfulness of an epistemic extension of the approach serves as a kind of indirect, defeasible *defense* of the relationship-based framework. In other words, the idea is that this feature of the approach seems like a mark in its favour, insofar as being able to account for a wide-range of phenomena with a simple idea is a mark in favour of that idea.

⁸ Thanks to an anonymous referee for pressing this worry.

4. Epistemic Blame and Epistemic Relationships

The core ideas of the relationship-based approach to moral blame extend naturally to the epistemic domain, enabling an understanding of epistemic blame in terms of a distinctive kind of relationship modification. In particular, we can understand judgments of epistemic blameworthiness as judgments (however implicit) that someone has done something that indicates they fall short of the normative ideal of some *epistemic* relationship they stand in with some other person. And actual epistemic blame responses can be understood as consisting in modifications to certain intentions and expectations one has towards that person, in a way made fitting by the judgment of epistemic blameworthiness.

On my approach, epistemic relationships are sets of intentions, expectations, and attitudes people have towards one another that are oriented towards their *epistemic agency* in distinctive ways. A promising way of approaching this idea is to focus on the fact that, whether we are aware of it or not, most of us have intentions and expectations towards one another simply in virtue of the fact that we are epistemic agents in an epistemic community. This is another way of putting the platitude that we rely on one another in a generic way as sources and distributors of information. We rely on one another for such basic things as finding out where the nearest grocery store is, or when the next bus arrives. Moreover, we don't *merely* rely on one another, for example, in the way a compulsive gambler might find themselves relying on a sports team in order to win a lot of money. Under certain conditions, at least, we also often have a degree of *confidence* in one another, in the sense that we expect (in both a normative and predictive sense of 'expect') the intellectual activity of others to meet certain epistemic criteria. Enjoying the status of an epistemic agent—of someone capable of concern for the epistemic justifiability of their actions and attitudes—puts one in a certain position vis-à-vis the lives of other inquirers: as an epistemic agent, one is a paradigm source of testimonial knowledge. As a paradigm source of testimonial knowledge, an epistemic agent's word ought to be epistemically trusted, unless a hearer has good reason not to. Members of an epistemic community stand in a relationship of mutual epistemic trust. I'll call this the 'general epistemic relationship'.

Epistemic trust is not trust simpliciter. Trusting someone entails, among other things, that you do not think they are out to deceive you or otherwise present you with purposefully misleading information. *Epistemically* trusting someone entails only that you have confidence that they are a reliable source of information—that believing what they say is a way of arriving at a favourable ratio of true to false beliefs, or knowledge, or understanding, etc.—and that you are willing to rely on them as such (McCraw 2015).⁹

Note that this way of expanding the idea of an epistemic relationship does not commit one to 'anti-reductionism' in the testimonial justification debate (Greco 2015). Both anti-reductionists and

⁹ A person can be untrustworthy on some matter but nevertheless be epistemically trustworthy on that same matter, and vice versa.

reductionists about testimonial justification can agree that, unless we have good reason not to, we should epistemically trust the word of other epistemic agents. Anti-reductionists can agree, because they maintain that testimony is a *sui generis* source of justification. Reductionists can agree, because they maintain that we often possess other sources of justification (abductive inference, knowledge of track records, etc.) in support of epistemically trusting the word of others.

Specific details about the nature of this relationship will need filling in. Who exactly counts as a member of the epistemic community? What counts as a good reason to suspend a presumption of epistemic trust? What kinds of circumstances might appropriately mitigate a suspension of epistemic trust, despite the fact someone has shown themselves to be unreliable? How finely grained should our presumptions of epistemic trust be individuated? That is to say, how far-reaching ought the implications be (regarding our epistemic trust) when other members of the epistemic community reveal themselves to be unreliable in a specific circumstance, or in a specific way? After all, people can be highly reliable on one subject matter, while being unreliable on another. These questions will be settled by our best substantive epistemological theories. They go hand in hand with our understanding of foundational epistemological issues such as the nature of justified belief, and the nature of testimonial justification. In the interest of remaining neutral on those substantive issues, I will stick with the following purposefully skeletal claim about the general epistemic relationship: the *normative ideal* of the general epistemic relationship specifies that each member intends to epistemically trust another unless one has good reason not to. And each member expects others to actually *be* epistemically trustworthy in this way, as well as that they will do the same of them. Even at this skeletal level, we can use the general epistemic relationship to make sense of at least two very basic ways in which we epistemically blame one another.

Epistemologists typically use examples of dogmatism (Brown forthcoming: 11), wishful thinking (Nottelman 2007: 75; Russel 2001: 40), hasty reasoning (McHugh 2012: 66), and certain kinds of biased cognition (see Bartolotti and Miyazono 2015) as examples of ways people can be appropriate targets of epistemic blame. What these examples have in common is that they are examples of intellectual conduct that tends to give others good reason to suspend their presumption of epistemic trust in would-be targets of epistemic blame, at least within some restricted domain, or on some specific matter. We can account for this with the following idea:

Epistemic Blame - Generic Type I: One party to the general epistemic relationship can appropriately epistemically blame another for conducting their intellectual lives in a way that gives one a good reason to suspend their presumption of epistemic trust in that person, at least within some restricted domain, or on some specific matter.

This thesis alone should not be taken to determine whether someone can appropriately epistemically blame another—there will be other factors to take into account (such as certain standing conditions, or other instrumental considerations). But setting that to one side, the sort of intellectual conduct specified here is constitutive of an impairment of the general epistemic relationship; it is one way of falling short of the normative ideal. By dogmatically holding a belief despite having excellent evidence to the contrary, for example, depending on the details of the situation, an agent may impair the general epistemic relationship in the sense that they fall short of at least one expectation that is constitutive of it, namely the expectation that others will actually *be* epistemically trustworthy, at least within some restricted domain, or on some specific matter. On my account, one’s modification to this expectation and corresponding intentions—for example, revising one’s intention to epistemically trust the word of another person on a subject matter they seem to be prone to think dogmatically about—in a way made fitting by the judgment that one’s epistemic relationship has been impaired, just is what it means to epistemically blame that person for their epistemic failing.¹⁰ Notice, once again, that this need not be a very weighty or strong kind of blame. Indeed, all of this can happen in pretty subtle ways. But that is exactly as it should be. When we epistemically blame someone for dogmatically believing something in the face of excellent evidence to the contrary, in many circumstances it need not be a particularly heated kind of blame.

This is an attractive idea because it enables us to make sense of numerous kinds of cases typically used in discussions of epistemic blame in the literature. But the account’s ability to classify cool responses as blame responses might seem to some like liability, as opposed to an advantage of the account. Consider: I observe the cashier in the supermarket doing something intellectually remiss while ringing my groceries through, and I modify my intention to epistemically trust them in response. Must I be understood as epistemically *blaming* the person? That might seem extreme. But it seems that this is what the relationship-based account classifies me as doing.

I am prepared to bite the bullet and insist that I epistemically blame the cashier. So long as I modify my intentions and expectations towards them, in a way made fitting by the judgment (however implicit) that they’ve impaired the general epistemic relationship, then I count as epistemically blaming them. Blame is a highly variable response on the relationship-based account. The fact that we can account for a wide variety of typical cases of epistemic blame with this very same idea is a reason to extend it, in a perhaps somewhat revisionary way, to this mundane supermarket case. I return to revisionary elements of the account in Section 4.

¹⁰ This feature of my account overlaps in interesting ways with recent work by Kauppinen (2018). Kauppinen puts ‘reductions in epistemic trust’ to work in a theory of ‘epistemic accountability’. Interestingly, he denies that there is such thing as epistemic blame.

The general epistemic relationship provides a helpful way of analyzing an even wider range of cases, ones that strike me as having something important to do with epistemic blame, even though they are not always discussed as such. Consider some of Miranda Fricker's (2007) cases of testimonial injustice. One of Fricker's central cases involves a man—Herbert—failing to epistemically trust the word of a woman—Marge—simply because of a prejudice Herbert has that women are in some sense overly 'emotional' and 'irrational' (2007: 9). According to Fricker, Herbert harms Marge in her 'capacity as a knower'. She claims that such harm is the distinctively *epistemic* harm of epistemic injustice. It seems Marge is entitled to epistemically blame Herbert (in addition to blaming him in other ways, such as morally).¹¹ We can account for this with the following idea:

Epistemic Blame - Generic Type II: One party to the general epistemic relationship can appropriately epistemically blame another for failing to epistemically trust their word (or the word of some third party) without good reason not to.

Such intellectual conduct is constitutive of an impairment of the general epistemic relationship; it is another way of falling short of the normative ideal. In our terms, when Herbert fails to epistemically trust Marge—simply on the basis of prejudicial stereotyping—this constitutes an impairment to the general epistemic relationship. Herbert reveals to Marge that he falls well short of the normative ideal of that relationship, which includes the intention to trust the word of another unless one has good reason not to. Marge's blaming of Herbert might consist in a judgment (however implicit) that Herbert has done something to impair the general epistemic relationship, and a modification to her intentions and expectations as a result. For example, minimally, Marge might modify her expectation that Herbert will trust her word unless he has good reason not to in the future. She might also modify her willingness to enter into other kinds of relationship with Herbert. Perhaps this will be accompanied by emotional responses such as indignation and resentment. The connection between epistemic blame and epistemic injustice has not, to my knowledge, been given a great deal of attention in the literature. The relationship-based account's ability to make headway here is an additional virtue of the account.

While this section falls short of presenting a fully developed relationship-based account of epistemic blame, it provides a sense of the promise of the idea that we can extend the relationship-based approach to the epistemic domain. With these details on hand, we can see in a more in-depth way how the relationship-based framework enables an explanation of the significance of epistemic blame without

¹¹ Fricker herself appears to deny this claim. However, she denies it because of the historical nature of her case. For those who agree with Fricker about that, I propose transposing this case to present day. See Povlarchy 2020 for discussion.

raising worries about whether talk of indignation, resentment, and other more heated kinds of response are out of place in the epistemic domain. Epistemic blame, like moral blame, is just another kind of relationship modification. It has a wide range of manifestations—from revisions to our expectations that someone will trust our word unless they have good reason not to in the future, to revisions to one's intention to epistemically trust the word of another person on a subject matter they seem to be prone to dogmatic thinking about. Such modifications have more significance (for you, and for the other person) than simply assigning someone a bad epistemic grade. But they are also familiar parts of many people's epistemic lives.

5. Objections to the Relationship-Based Framework

I have addressed a couple of worries already. In this section I respond to three further potential objections.

i) It seems there can be cases in which a person blames someone, but does not *modify* their intentions and expectations towards that person in any way. Consider a husband who fails to bring his partner anniversary flowers for the 7th year in a row. His partner might blame him for this, but not make any *modifications* to his intentions and expectations. It's just too familiar by now. In response, I do not find it implausible that, even in a case like this, one may modify their intentions and expectations ever so slightly, in a way made fitting by the judgment that the relationship falls short of some normative ideal. Given the contextual features of the case, even very subtle adjustments may provide a plausible way of understanding what someone's blame-response consists in. But we can also make a minor adjustment to the way I've presented the relationship-based account. In addition to making modifications to intentions and expectations central to blame, we can add the plausible idea that reaffirmations of certain intentions and expectations are equally central. Thus, it may be true that the man blames his husband in the above example without making any modifications to his's intentions and expectations. But it's plausible that when the man judges his husband blameworthy, he nevertheless reaffirms, or in some way maintains the intentions and expectations he originally formed in response to the impairment a number of years ago.

ii) Angela Smith (2013: 137) has argued in her own way that Scanlon is forced to classify certain kinds of relationship modification as instances of blame that intuitively are not instances of blame (this is a variation on the objection, considered earlier, that the account misclassifies the cashier case). She uses the example of a mother intending to care for and appreciate her son even *more* upon learning of a particular crime he has committed. It is important to remember that there is a central role for judgments of *blameworthiness* in the relationship-based account, as well as the claim that, in order to count as blame responses, modifications to one's relationship with another person must be ones that are made fitting by that judgment. It is implausible that the mother meets the latter condition. That is, it doesn't seem right to say that her response is one that is made fitting by the judgment that her son is *blameworthy*, at least not

from the point of view of morality. Indeed, we might say that her feelings of partiality towards her son are, in some sense, getting in the way of her ability to blame her son. Moreover, technically speaking, the account does not merely maintain that modifications to one's expectations and intentions need to be made fitting by a judgment that X is *blameworthy*. Rather, it maintains that these modifications need to be made fitting by a judgment that X has done something to *impair one's relationship with X*. In Smith's mother-son case, the mother simply judges that the son is blameworthy, perhaps in the sense that he has done something to impair *a relationship*, as to opposed to their own parent-child relationship. So, her reaction need not be classified as a blame response according to the account.

Still, it seems fair to request an explanation of which modifications to our relationships count as *fitting* responses to judgments of relationship impairment, and why that is the case. Some might argue that this talk of 'fittingness' is rather vague, and potentially ad hoc. The account ultimately relies, in part, on a shared grasp of the sorts of things we typically associate with blame and relationship impairment. In this regard, it is no different from the most sophisticated competing account of epistemic blame in the literature, due to Jessica Brown (forthcoming). Brown has drawn on George Sher's (2006) account of moral blame, arguing that it extends naturally to the epistemic domain. According to Sher, blame is a *characteristic* set of dispositions to feel and behave, unified by a particular belief-desire pair. The sorts of dispositions Sher has in mind are just those we 'ordinarily associate with blame' (Sher 2006: 93-114). Neither Sher nor Scanlon are in the business of giving a reductive analysis of blame. The primary contributions of both of their accounts are their explanations of what *unifies* a diverse range of phenomena in our practice of responding to one another for moral failings. In Sher's case, it is a belief-desire pair. In Scanlon's, it is the modification of intentions and expectations in response to the judgment that one's relationship has been impaired.

I argued above that we can also embrace a somewhat revisionary approach to the concept of blame (Scanlon 2008: 211). Perhaps some things that get counted as blame on our account do not seem intuitively quite like blame; this may simply be a cost proponents of the relationship-based framework should be willing to take on, in light of theoretical advantages the account has otherwise. In such cases, it is open to for one to admit that perhaps we don't think of them quite fully as blame-responses, but argue that we should. Scanlon himself has done much to highlight the theoretical advantages of his own account (he explores issues such as moral luck, and the relationship between blame and control) (Scanlon 2008: 123-213). This paper is in part an exploration of a further theoretical advantage, namely the applicability of its central ideas to the epistemic domain.

iii) It is widely regarded as an important desideratum for any theory of blame that it can account for the possibility of self-blame. It may seem unclear how an account of blame that grounds our understanding of blame in our *relationships* with one another can account for the intuitive idea that one

can blame *oneself* for an action or attitude. To take self-blame on board, we need to countenance the idea one can have a relationship with oneself, and that there are ways we can impair this relationship, as well as modify it in response to judgments of impairment. I suggest, along with Scanlon (2008: 155) that one *does* has a relationship with oneself, perhaps even in the ordinary sense of that term. We tend to rely on ourselves, believe in ourselves, know ourselves, hate ourselves, love ourselves, etc. These are all things that arguably figure essentially in our relationships with *others*, things that make those relationships the sorts of things that they are.

This claim is even more plausible when we recall the technical sense in which I am using the term ‘relationship’. In the framework, roughly speaking, relationships are just sets of intentions and expectations people have towards one another. To accommodate the idea of a relationship with oneself, we simply have to allow that there can be certain intentions and expectations that one can have towards oneself—intentions and expectations concerning how one will act toward oneself, and concerning the feelings one has for oneself. This is a familiar idea. When one blames oneself, on this picture, one judges that one falls short of some normative ideal (one let oneself down, say) and modifies one’s understanding of one’s relationship with oneself in a way made fitting by that judgment (for example, perhaps one places less trust in one’s ability to follow through in the future, and perhaps this is accompanied by guilt or anger, etc.).

6. Other Sources of Skepticism about Epistemic Blame

There remains a bigger picture issue regarding my defense of the idea that there is a distinctively epistemic kind of blame. I have only argued that my account speaks to, and ultimately deflates, *one* reason to be skeptical of epistemic blame. Namely, I have argued that it avoids the problem of cool judgment. In this final section, I address a few other possible sources of skepticism about epistemic blame, and argue that that they either fall apart on their own terms, or can be deflated with the help of the relationship-based account of epistemic blame.

6.1 Epistemology is a Branch of Ethics

Some authors are skeptical about epistemic blame because of their views about the relationship between ethics and epistemology.

There are a wide variety of attitudes towards this relationship. For example, some recent virtue epistemology casts doubt on the idea that there is a firm distinction to be drawn between ethics and epistemology. In their detailed examination of the intellectual virtues, Roberts and Wood say something to this effect:

We find it unhelpful to try to draw a strict line between the intellectual and the moral virtues. So we will speak of intellectual humility, intellectual courage, intellectual generosity, where more traditional usage might speak of a moral virtue applied to an intellectual context [...]The difference between our study and a study in virtue ethics is simply that we are interested in the relations between the virtues and the intellectual goods” (Roberts and Wood 2007: 60).

And consider the following from Linda Zagzebski:

I will argue that the intellectual virtues are so similar to the moral virtues . . . that they ought not to be treated as two different kinds of virtue. Intellectual virtues are, in fact, forms of moral virtue. It follows that intellectual virtue is properly the object of study of moral philosophy . . . It will take most of this book to demonstrate that epistemic evaluation is a form of moral evaluation . . .’ (Zagzebski 1996: xiv, 6)

While these philosophers are all primarily concerned with the intellectual virtues, as opposed to claims about epistemic blame, their views about the relationship between moral and intellectual matters are potentially relevant to a study of epistemic blame. If one can argue that there is little use—perhaps even something mistaken—in distinguishing between moral and epistemic *virtue*, one might argue something similar about epistemic blame. Indeed, on a more general note, Zagzebski herself claims that ‘[e]pistemic evaluation just is a form of moral evaluation, [...] it follows that normative epistemology is a branch of ethics’ (Zagzebski 1996: 258).¹²

These are controversial claims. It is perhaps unclear how worried a theorist of epistemic blame should be about a form of skepticism about epistemic blame that rests on the assumption that epistemology is a ‘branch of ethics’. However, it would be helpful if we could say something positive to deflate this source of skepticism. One response is to start with the observation made in Section 2 that not all theories of blame seem as though they would naturally extend to the epistemic domain. As I argued there, not all theories of blame seem as though they would be able to provide compelling analyses of the way we respond to one another for our epistemic failings, when what we are doing is something beyond mere negative epistemic evaluation. The fact that such a theory does not naturally extend to the epistemic

¹² Zagzebski connects this latter claim to blame directly in a footnote: ‘after writing these words I discovered that H.H. Price (1954) used this same worry to explain his reticence about attaching moral concepts to cognitive activity: “It seems to me that we are all far too much addicted to blaming people as it is. If we are to be allowed, or even encouraged, to blame them for the way they direct their thoughts, as well as for their actions, there will be a perfect orgy of moral indignation and condemnation, and charity will almost disappear from the world” (quoted in Dearden, Hirst, and Peters, 1972, 364-5)’ (Zagzebski 1996: 258).

domain should not *ipso facto* discount it as a theory of blame. Perhaps it is a good theory of moral blame, even if it doesn't naturally extend to the epistemic domain. This suggests that the kind of evaluations we make in the epistemic domain really can come apart in different ways from evaluations in the moral domain.¹³ For this reason, it seems odd to say that epistemology is just a branch of ethics. So, we should take seriously the possibility that there are different kinds of blame.

6.2 Ontological Parsimony

In a broadly related vein, but focusing more specifically on the present topic, some have argued for the 'reduction' of epistemic blame to moral blame. More specifically, some have argued that, for any epistemic failing, either something of moral or prudential significance hangs on the failing, or it does not. If it does, then the kind of blame, or attribution of responsibility, that properly attaches to the failing is moral or prudential in nature, not epistemic. If nothing of significance hangs on the failing, then blame cannot be the appropriate response to the failing (Papineau 2012, 4). Thus, all instances of epistemic blame, epistemic blameworthiness, and epistemic irresponsibility just *are* instances of other practical or moral kinds of blame, blameworthiness, and irresponsibility. The basic motivation for this view is ontological parsimony. If we can account for the appearance of distinctively epistemic kinds of blame, blameworthiness, and irresponsibility without positing additional, *sui generis* forms of these things, we should.

I find this line of argument unconvincing. First, why must blame for a failing that has moral or prudential significance itself be either a moral or prudential kind of blame? Presumably the thought is that, if X-ing is either *morally* significant or *prudentially* significant, the relevant blame response towards someone for X-ing should somehow appropriately reflect this. That is, we must respond to something of moral significance in a moral way—not, for example, an epistemic way. Regardless of whether that's true, one thing it is important to note is that there are other kinds of significance beyond moral and prudential significance. For example, perhaps there is a kind of epistemic significance. An example of something of epistemic significance might be the discovery that $E=MC^2$, or the discovery of whatever it was that caused so many dinosaurs to disappear. Such discoveries may also prove to be of prudential and

¹³ Zagzebski might respond that this is not the only result suggested by my observation. Instead, perhaps my observation suggests that the 'epistemic' branch of moral philosophy has its own kind of blame. Or perhaps my observation suggests that there is no such thing as epistemic blame (because we don't blame each other in the epistemic branch of morality, for example). Regarding the first alternative, if we go this route, then it seems to me that any disagreement between people like Zagzebski and myself is a merely verbal one. Regarding the second route, I of course disagree; after all, I have just presented an independently motivated framework for thinking about what epistemic blame might be.

moral significance. But there is something intuitive about the idea that they are significant from a purely epistemic point of view, too.¹⁴

If there is such thing as epistemic significance, then a key premise of this argument is false. It does not follow that, just because something of moral or prudential significance hangs on someone's epistemic failing, that they can only appropriately be blamed in a moral or prudential way. Something of epistemic significance might *also* hang on their epistemic failure. And they might be blameworthy in an epistemic way for that failing, too. Of course, it is notoriously difficult to explain intuitions about 'epistemic significance' without bottoming out on claims about moral, or practical, or some other non-epistemic kind of value. I cannot make a decisive contribution to that issue here. However, the relationship-based account has resources to do something close enough for present purposes.

Perhaps it's true that no one particular truth, or item of knowledge has any more epistemic significance in its own right than any other. In my view, that is neither here nor there. Understanding the appropriateness of a distinctively epistemic kind of blame does not require appealing to some kind of impersonal or abstract form of epistemic significance. It only requires that one person's epistemic failing can be of epistemic significance *for* someone else, or themselves. And the idea of epistemic relationship impairment can help with just that. One person's epistemic failing can have epistemic significance for another, and for themselves, insofar as it impairs their epistemic relationship with that person. Even a case as mundane as the grocery store example from earlier involves a failing of some epistemic significance for the epistemic agents involved. The clerk's failing makes it fitting for me to modify certain intentions and expectations I have towards them, intentions and expectations that are oriented towards their epistemic agency in distinctive ways. The idea that this episode has epistemic significance for me fits naturally with the more general idea that when we blame people, we typically do so in part because their failing has significance for us (or others close to us).

Of course, the reductionist may reject the entire notion of an epistemic relationship altogether, perhaps on grounds of ontological parsimony. But ontological parsimony doesn't come for free, and I would respond by pointing out the usefulness of this notion. After all, as the present discussion indicates, we can get considerable theoretical mileage out of the idea. I leave for another time an exploration of other issues and debates we might be able to illuminate with the concept of an epistemic relationship.

6.3 Blame and Control

Finally, another potentially salient difference between the epistemic and moral domains centers on the notion of control. Consider some well-known observations about the differences in the kind of control we

¹⁴ To put it one way, surely coming to know that $E=MC^2$ is more significant than coming to know how many blades of grass are on your lawn, regardless of what then happens from a moral or practical point of view.

have over our doxastic states in comparison with the kind we have over our actions. Imagine being offered \$1,000,000 to raise your arm, or to flip a light switch. Under normal circumstances you'd be able to earn this desirable sum of money. But now imagine trying to form a belief that Clinton became president in 2016, just because you are being offered \$1,000,000. Even when such a reward is on offer, one cannot simply believe that *p* so long as it goes against one's evidence.¹⁵

Many are impressed by the idea that it would be inappropriate to blame someone for something that was outside their control. Showing that something was outside your control is a paradigm *excuse* defense. If our doxastic states are outside our control, how could we appropriately be blamed for them? Given that doxastic states are supposedly paradigm targets of epistemic blame, this may cast doubt on the idea that there is a distinctively epistemic kind of blame.

This is a longstanding issue, and there are a wide range of things to say. For a start, there are many views about the nature of doxastic control. Most agree that we do not have 'basic voluntary control' over our beliefs (Alston 1988). But perhaps we have a kind of indirect control over our doxastic states, such as by making decisions about who to associate with, or what bodies of evidence to consider (Alston 1988; Rettler 2017). Perhaps doxastic control simply amounts to a kind of reasons-responsiveness (McHugh 2015; Mitova 2011; Owens 2000; Ryan 2003; Hieronymi 2009; Steup 2008; Shah 2002).

But we can also challenge the more basic assumption that blame implies control. For example, some have appealed to the fact that there seem to be a number of areas, outside epistemology, where blame appropriately applies in the absence of control (Ryan 2005); others have argued that while blame implies some kind of *ability*, this falls short of the claim that blame implies control (Chuard and Southwood 2009).

As with the problem of cool judgment, my preferred take on this issue is to note that the connections between blame and control force us to get clear, among other things, on what sort of thing blame is. If blame is one sort of thing—say, for example, a manifestation of emotional responses like indignation and resentment—then perhaps it really is hard to see how it could be appropriate in the absence of control. But if blame is something else—for example, if it can consist in certain kinds of relationship modification, as I've argued—then perhaps it's not. Indeed, in my view, when we ground our understanding of blame in terms of our relationships with one another, it is not hard to see how blame can be appropriate in the absence of control. Consider the fact that, often times, attitudes and actions that we do not (seem to) have voluntary control over can nevertheless express a great deal about how we regard one another. They often imply a great deal about our attitudes, intentions, and expectations towards one

¹⁵ And it is not merely a matter of going *against* one's evidence. McHugh argues that we are incapable of forming beliefs for which we take ourselves to have no supporting evidence, or even insufficient evidence (McHugh 2012: 69).

another. As Angela Smith has pointed out, sometimes things we do over which we do not have direct voluntary control say *more* about who we are, and the regard we have for others, than the things we do have control over. Smith discusses a variety of responses that she calls ‘involuntary responses’, which, she argues, are an important indication of someone’s underlying values and ‘moral commitments’—of ‘who they are’, so to speak. Her interest in these responses is that, ‘As a number of philosophers have pointed out, their alleged “passivity,” far from undermining their attributability to persons, may actually be the strongest mark of their genuineness and sincerity’ (Smith 2005: 242). The sorts of responses Smith has in mind include things like *noticing*, *neglecting*, ‘*what occurs to us*’, and a variety of other ‘involuntary *reactions*’, such as jealousy, contempt, and regret.

A person’s beliefs and other doxastic states fit comfortably on this list of ‘involuntary responses’. What we are disposed to believe, under particular circumstances (i.e. given our evidence, and background beliefs, and so on), often says much about who we are, at least from the perspective of our epistemic agency. If I’m disposed to believe things about Prince Harry’s love life simply on the basis of some trashy tabloids I read, this says something about my ‘intellectual character’. Indeed, there is a wide range of vocabulary—familiar from virtue epistemology—that suggests as much. We often refer to each other as ‘open-minded’, ‘gullible’, ‘conscientious’, ‘curious’, ‘intellectually autonomous’, and so on. Importantly, the doxastic dispositions implied by such terms express features of our deeper evaluative commitments, and they do so irrespectively of any voluntary control we may or may not have over them.

Adjustments to our expectations and intentions towards one another, in light of our epistemic conduct, can be apt because of what a person’s epistemic conduct and doxastic states say about a person—about their intellectual character—and what this in turn *means* for you, insofar as you stand in certain relationships with that person (including, of course, the general epistemic relationship). Adjustments to our expectations and intentions towards one another—the kind that count as a sort of relationship modification made fitting by a judgment that one’s epistemic relationship with another falls short of the normative ideal of that relationship—is what epistemic blame *consists in* on my account. So epistemic blame can be an appropriate response to a person for being in certain doxastic states in the circumstances, even if (or maybe even partially *because*) these are attitudes over which that person has no immediate voluntary control.¹⁶ Thus, the fact that belief is a paradigm example of a state over which agents do not have immediate voluntary control provides no reason to worry about whether epistemic blame is, generally, a legitimate notion. The relationship-based account of epistemic blame has significant

¹⁶ Of course, there are important limits to this point. Not everything I do, over which I lack voluntary control, is something I’m open to blame for. But this fits naturally with the obvious point that, not everything I do is open to interpretation as a way of *impairing* my relationships with other people.

resources for handling, and, in my view, ultimately deflating worries about the very idea of epistemic blame that stem from the putative connection between blame and control.

References

- Bartolotti, L. & Miyazono, K. 2015. *Erkenntnis*, 81: 275-296.
- Brown, J. forthcoming. What is epistemic blame? *Noûs*.
- _____. 2017. Blame and wrongdoing. *Episteme*, 14:3.
- Chuard, P. & Southwood, N. 2009. Epistemic norms without voluntary control. *Noûs* 43:4.
- Descartes, R. 1955. Meditation 4. In *Philosophical Works of Descartes, Vol. I*, Haldane and Ross (eds.), reprint, Dover: New York.
- Feldman, R. 2004. The ethics of belief. In E. Conee and R. Feldman, *Evidentialism*. OUP: Oxford.
- Fricke, M. 2007. *Epistemic Injustice*. OUP: Oxford.
- _____. 2016. What's the point of blame? *Noûs*, 50: 165-183.
- Gokhale, S. 2019. Who needs blame?: answerability without expressed blame. *CUNY Academic Works*.
- Hieronymi, P. 2004. The force and fairness of blame. *Philosophical Perspectives*, 18:115-48.
- Kauppinen, A. 2018. Epistemic norms and epistemic accountability. *Philosophers' Imprint*.
- Littlejohn, C. Forthcoming. A plea for epistemic excuses. In *The New Evil Demon: New Essays on Knowledge, Justification and Rationality*, Dutant, J. (ed.), OUP. Oxford.
- McCraw, B. 2015. The nature of epistemic trust. *Social Epistemology* 29: 413-430.
- McHugh, C. 2012. Epistemic deontology and voluntariness. *Erkenntnis* 77.
- _____. 2015. Attitudinal control. *Synthese*, 194:8.
- Menges, L. 2017. The emotion account of blame. *Philosophical Studies* 174.
- Nottelmann, N. 2007. *Blameworthy Belief: A Study in Epistemic Deontologism*. Springer: Dordrecht.
- Owens, D. 2000. *Reason Without Freedom: The Problem of Epistemic Normativity*. Routledge.
- Papineau, D. 2012. There are no norms of belief. In *The Aim of Belief*, Chan, T. (ed.). OUP: Oxford
- Peels, R. 2016. *Responsible Belief: A Theory in Ethics and Epistemology*. OUP: Oxford.
- Plantinga, A. 1993. *Warrant: The Current Debate*. OUP: Oxford.
- Poviarchy, A. 2020. Responsibility for testimonial injustice. *Philosophical Studies*. Online first.
- Rettler, L. 2017. In defense of doxastic blame. *Synthese*, online first.
- Russel, B. 2001. Epistemic and moral duty. In Steup, M. (ed.) *Knowledge, Truth, and Duty. Essays on Epistemic Justification, Responsibility, and Virtue*, OUP: Oxford.
- Ryan, S. 2003. Doxastic compatibilism and the ethics of belief. *Philosophical Studies* 114.
- Scanlon, T.M. 1998. *What We Owe to Each Other*. Harvard University Press: Cambridge.

- _____. 2008. *Moral dimensions: Permissibility, Meaning, Blame*. Belknap Press: Cambridge, MA.
- Sher, G. 2006. *In Praise of Blame*, OUP: Oxford.
- Smith, A. 2013. Moral blame and moral protest. In *Blame: Its Nature and Norms*, Coates, D.J & Tognazzini, N.A (eds.). OUP: Oxford.
- _____. 2005. Responsibility for attitudes: activity and passivity in mental life. *Ethics* 115: 236-271.
- Steup, M. 1999. A defense of internalism, in Pojman, L. (ed.), *Theory of Knowledge: Classical and Contemporary Readings*. Belmont, CA: Wadsworth.
- Strawson, P. 1962. Freedom and resentment. In *Proceedings of the British Academy, Vol. 48: 1962*, G. Watson (ed.). OUP: Oxford.
- Van Woudenberg, R. 2009. Responsible belief and our social institutions. *Philosophy*, 84:1.
- Wallace, R.J. 1994. *Responsibility and the Moral Sentiments*. Harvard University Press.
- _____. 2013. Dispassionate opprobrium. In *Blame: Its Nature and Norms*, Coates, D.J & Tognazzini, N.A (eds.). OUP: Oxford.
- Williamson, T. forthcoming. Justifications, excuses, and sceptical scenarios. In *The New Evil Demon: New Essays on Knowledge, Justification and Rationality*, Dutant, J. (ed.), OUP. Oxford.
- Wolf, S. 2011. Blame, Italian style. In *Reasons and Recognition*.