

Scepticism about Epistemic Blame*

Tim Smartt

timjsmartt@gmail.com

Forthcoming in *Philosophical Studies*

Abstract

I advocate scepticism about epistemic blame; the view that we have good reason to think there is no distinctively epistemic form of blame. Epistemologists often find it useful to draw a distinction between blameless and blameworthy norm violation. In recent years, this has led several writers to develop theories of ‘epistemic blame.’ I present two challenges against the very idea of epistemic blame. First, everything that is supposedly done by epistemic blame is done by epistemic evaluation, at least according to a prominent view about the social role of epistemic evaluation. Parsimony considerations count against introducing an idle mechanism that does the same work as an existing one. Second, no current theory of epistemic blame includes a plausible account of the force of epistemic blame or the practices that could express it. I conclude that we should give up the notion of epistemic blame.

1 Introduction

A common strategy for defending highly demanding epistemic norms is to claim that sometimes an agent’s failure to satisfy a norm is blameless. Drawing a distinction between blameless and blameworthy norm violation provides a powerful defence of demanding norms—it allows one to maintain that they are

*Earlier versions of this paper were presented at Brown University, the University of Colorado Boulder, the Southampton-Humboldt Normativity Conference and the Notre Dame/Northwestern Graduate Epistemology Conference. I’m grateful to these audiences, as well as to Brian Hedden, Mark Colyvan, Hannah Tierney, Adam Piovarchy, Amy Flowerree, Brandon Price, Patrick McKee and Jonathan Weid for very helpful comments that improved this paper.

genuine norms, while avoiding the objection that they're *excessively* demanding by creating a way to temper our judgments of agents who fail to satisfy them. For instance, Bayesians standardly claim that an agent who fails to be logically omniscient, or update by conditionalization, or have probabilistically coherent credences, might nevertheless be blameless for some of these failings (Christensen 2004).

In recent years, the most prominent example of this strategy can be found in the Williamsonian tradition, where writers commonly rely on the blameless/blameworthy distinction to defend highly demanding knowledge norms about various issues. Consider Williamson's (2000, 47) defence of the knowledge norm of belief, according to which one should only believe what one knows. On this view, it's always wrong to believe p when one does not know p .¹ Nevertheless, Williamson (forthcoming) holds that one can blamelessly violate the norm—say, if one's belief was formed by impeccable cognitive dispositions, or if one believes p and has good evidence that one knows p (but in fact one doesn't know p). So, in the 'new evil demon problem'—involving a brain in vat whose experiences and cognitive dispositions are duplicates of your own—the brain in a vat who falsely believes that it has hands blamelessly violates the knowledge norm of belief.²

This strategy generates a puzzle. Moral blame is familiar—it's a commonplace feature of our interpersonal lives and it's been carefully studied by philosophers, social psychologists, and others. But the notion of 'epistemic blame' is much less familiar. It's natural to wonder whether there is a special kind of blame which targets epistemic, rather than moral, failings.

In this paper I advocate scepticism about epistemic blame; the view that we

¹See Williamson (2005, 2013, 2014) for similar arguments in defence of knowledge norms of practical reason, assertion, and justification. Knowledge-firsters often make similar moves. For example, this type of argument is used by Hawthorne and Srinivasan (2013, 20) to defend the knowledge norm of disagreement. On their view, whether one ought to be 'conciliatory' or 'steadfast' in a case of peer disagreement depends on which option will result in knowledge; S should be conciliatory, steadfast, or suspend judgment regarding p , if and only if doing so results in S knowing p . But they hold that one can blamelessly violate the norm if one opts for what seems like the knowledge-conducive strategy, but in fact isn't. If you and your friend arrive at different calculations about how to split the dinner bill, and you've been drinking whereas she hasn't, Hawthorne and Srinivasan take it that you could blamelessly be conciliatory—even if being steadfast was the knowledge-conducive option in this case.

²In general, these arguments rely on a normative structure in which primary norms—which identify what one fundamentally ought to do—are naturally connected with derivative norms—which identify what one usually ought to do to satisfy the primary norm. This structure produces judgments of blamelessness with a systematic elegance; when an agent satisfies a derivative norm, they are often blameless for failing to satisfy a corresponding primary norm. See, Lasonen-Aarnio (2010, 12-15), Hawthorne and Stanley (2008, 586) and Williamson (forthcoming).

have good reason to think there is no distinctively epistemic form of blame. I provide two arguments for this view, which draw on a common source. A hallmark of moral blame is that it carries a special force, which allows it to rise above the level of mere evaluation. When we blame someone, we do more than simply register that they've violated a norm. Pamela Hieronymi (2004, 116) captures this feature nicely:

Blame, it is thought, goes beyond simple description or mere grading...Being morally blamed involves a more serious sort of criticism than being told your vocal performance was flat, your cooking bland, your conversation dull, or your sentences opaque. Blame, unlike mere description, carries a characteristic depth, force or sting.

This feature is so central that a theory of blame is deficient if it fails to adequately account for its special force. For instance, Gary Watson (2004, 226) claims that when we blame, we're not merely 'moral clerks, recording moral faults' from a detached standpoint. Similarly, Scanlon (2008, 127) claims that blame has distinctive weight and any account that reduces it to 'a pointless assignment of moral grades' is unacceptable.

This central feature provides a simple test for whether a type of criticism or reaction is a species of blame: does it carry a special force that allows it to rise above the level of evaluation? I argue that the notion of epistemic blame fails this test: all the extant accounts of epistemic blame fail to exhibit a feature that's necessary for a kind of negative reaction to be a species of blame. In [Section 2](#), I clarify the notion of epistemic blame and summarise several accounts of it that have recently been developed. In [Section 3](#) and [Section 4](#) I argue that all these accounts fail my test. In [Section 3](#), I argue that one account of epistemic blame—the relationship modification account—is not sufficiently different from mere epistemic evaluation. In [Section 4](#), I argue that our typical reactions to epistemic norm violations do not have the characteristic depth, force or sting needed to rise to the level of blame. Although several accounts attempt to capture this feature—especially emotion-based and belief-desire accounts—I argue that in doing so they provide an implausible picture of our normal practices of epistemic criticism. I conclude that we should give up the notion of epistemic blame. My view has two upshots. First, it uncovers an interesting fault line between epistemic and moral normativity: we have good reason to think that there's no purely epistemic kind of blame. Second—and more importantly—it rules out a common strategy for defending highly

demanding epistemic norms. If there's no epistemic blame, classifying epistemic failings as either blameworthy or blameless is inapt.

2 What is epistemic blame?

Before I present my arguments against epistemic blame, it's worth quickly clarifying the idea and reviewing how epistemologists have come to think about it. Epistemologists who accept the notion of epistemic blame claim that sometimes an agent is blameworthy for their *purely* epistemic failings. Just like we might blame someone if they violate a moral norm, those who accept epistemic blame claim that sometimes we might blame someone if they violate an epistemic norm.

Epistemic blame is supposed to be a special kind of blame, both distinct from, and continuous with, moral blame in certain respects. It's distinct from moral blame, since it targets epistemic failings, rather than moral failings that might be associated with our attitudes.³ For instance, an agent might be *morally* blameworthy for having a sexist belief, but not epistemically blameworthy,⁴ whereas an agent might be *epistemically* blameworthy for believing against the evidence, but not morally blameworthy. Epistemic blame is also supposed to share some of the central features of moral blame; it's supposed to be a more severe type of response than mere evaluation, one which carries a characteristic force and is a useful mechanism for encouraging norm uptake.

³Whether one accepts that agents can be *morally* blameworthy for their attitudes might depend on whether one accepts 'doxastic wronging'—the view that an agent's unexpressed attitudes can violate moral norms (Basu 2018, 2019; Basu and Schroeder 2019). If this view is true, then agents can be morally blameworthy for their attitudes. But this is distinct from the question of whether agents are epistemically blameworthy for their attitudes. It's also distinct from a view quite similar to doxastic wronging that holds that although beliefs *themselves* don't wrong others, they can cause one to perform an action that wrongs someone, or even just unacceptably runs the risk of wronging someone. Just as we might think that some actions that don't actually harm anyone are moral failings because they easily could have—for example, driving drunk or dumping carcinogens in a river—one might think that some unexpressed attitudes—for example, sexist or racist attitudes—are moral failings because they carry a significant risk of causing downstream harm. Nevertheless, the species of blame at work here is still moral blame. Whether an attitude constitutes a wrongdoing, causes a wrongdoing, or recklessly risks causing a wrongdoing, we're still reacting to a *moral* failing.

⁴Several writers accept a view along these lines regarding profiling, generalisations, and other kinds of statistical inferences about individuals based on their membership in social groups. Often our epistemological and moral evaluations of these inferences are in tension. See, for example, Gendler (2011, 57).

2.1 Three accounts of epistemic blame

Several writers have recently developed theories of epistemic blame.⁵ These can be grouped into three kinds of accounts: emotion-based accounts, belief-desire accounts, and relationship modification accounts.

According to **emotion-based accounts**, epistemic blame consists in a type of negative emotional response to perceiving that an agent has violated an epistemic norm. The set of emotions are familiar from how we typically respond to people's moral failings: anger, indignation, resentment, and so on. These accounts stress that our negative reactions to others' purely epistemic failings centrally involves an emotional component: we *feel* angry, upset, frustrated, and so on.⁶

According to **belief-desire accounts**, epistemic blame consists in a characteristic set of dispositions that are connected to a belief-desire pair. This view is due to Jessica Brown (2017, 2020a, 2020b), and is based on George Sher's (2006) account of moral blame. Sher takes it that moral blame consists in a characteristic set of dispositions—for instance, dispositions to remonstrate, feel upset, request justifications or apologies—that are causally connected to a belief-desire pair; namely, the *belief* that the targeted agent violated a moral norm and the *desire* that they hadn't. Brown holds onto the idea that blame involves a characteristic set of dispositions to feel and behave in ways that are causally connected to a belief-desire pair, but changes the norm violation that features in the belief-desire pair to epistemic norm violation. On her view, in cases of epistemic blame we believe that an agent violated an epistemic norm and we desire that they hadn't. This gives rise to dispositions to react in characteristic blame-like ways, such as privately feeling upset with them or publicly rebuking them by saying things like '*What do you mean the earth is flat?!*'

According to **relationship modification accounts**, epistemic blame consists in modifying one's interpersonal relationship with an agent in response to perceiving that they have violated an epistemic norm (and expecting them to not violate such norms is part of the relationship).⁷ In a nutshell, epistemic blame consists in lowering our confidence in someone in response to what we per-

⁵Although theories of epistemic blame have only been developed quite recently, precursors can be found in the history of epistemology. For example, Tollefsen (2017, 357) highlights that Thomas Reid (1785/2002) wrote that 'epistemic ridicule' was an appropriate emotional reaction to someone whose beliefs blatantly disregard common sense.

⁶Tollefsen (2017), Rettler (2018), McHugh (2012), Nottelmann (2007), and Peels (2013).

⁷This view has primarily been defended by Cameron Boulton (2020, 2021a, 2021b). It's also accepted in Schmidt (2021, 2022) and discussed in Greco (forthcoming).

ceive as an epistemic failing. These accounts adapt Scanlon's (2008, 2013) account of moral blame, which holds that blame centrally involves adjusting our intentions, expectations, and attitudes towards those we blame; we might intend to have less to do with them in the future, or withdraw our affection for them. Relationship modification accounts of epistemic blame claim that the same dynamics are at work in the epistemological domain.

A crucial element of these accounts is describing the epistemic analogue of the relationship being modified. They tend to pick out trust as the main feature of our relationships that is modified by epistemic blame. For instance, Daniel Greco (forthcoming) provides the following as an example of epistemic blame:

Suppose I regard you as generally credible and reliable, and will take your word when you make claims about matters where I lack direct evidence. When you told me that the available evidence suggests that taking zinc supplements reduces the duration of the common cold, I believed you. Now, I learn that you've taken up belief in homeopathy. Plausibly, this should lead me to modify my epistemic relationship with you; I should no longer be as willing to trust that you are a reliable source of information; maybe you're worse than I thought at finding credible sources, or worse than I thought at interpreting the evidence you get from those sources—perhaps I'll rethink those zinc supplements. This modification may be a matter of degree—I may still believe you in certain cases (e.g., concerning matters where there is little room for interpretation), or will give your word *some* non-zero weight even concerning trickier questions, but our epistemic relationship will have been significantly altered.

Similarly, Cameron Boult (2020, 2021a, 2021b) uses the example of trust to show that epistemic relationship modification can be presented in the Scanlonian framework; paradigmatically, we epistemically blame another in response to their epistemic failing by intending to trust them less (or on a smaller range of issues) and we expect them to be less reliable sources of information.

3 The social role of epistemic evaluation

In this section I argue against the relationship modification account. One part of my test outlined in [Section 1](#) is that any species of blame must be a more

severe kind of response than evaluation; epistemic blame must have a different job description to plain old epistemic evaluation. I argue that the relationship modification account fails to deliver this.

3.1 Dogramaci on epistemic evaluation

Epistemic evaluation plays a role in our social life. Several epistemologists have recently investigated what human needs are served by practices of epistemic evaluation. According to one plausible way of thinking about the social role of epistemic evaluation, everything that's supposedly done by epistemic blame—on the relationship modification account—is done by epistemic evaluation. The account of evaluation is mainly due to Sinan Dogramaci (2012, 2015), but is part of a broader tradition investigating the social role of evaluative epistemological notions due to Edward Craig (1990).⁸ Craig's original view is, 'to put it briefly and roughly, the concept of knowledge is used to flag approved sources of information.' (Craig 1990, 11) Other contributors to this program have argued that it's not only knowledge attributions that have this role; evaluative notions such as justification, rationality, and the like, have similar social functions. Let's look at Dogramaci's view as an example of this way of thinking about the practical purposes of epistemic evaluation.

Dogramaci argues that the role and utility of ordinary epistemic evaluations—such as knowledge attributions or judgments about rationality—is to foster interpersonal coordination within an epistemic community. These practices allow a community to be set up in such a way that true beliefs spread through the community. The primary way they do so is by establishing a system through which members of a community become one another's 'epistemic surrogates', whereby each agent's beliefs and experiences expand the pool of evidence available to other members of the community for acquiring true beliefs. Epistemic evaluations allow testimony to become a safe and efficient way to promote true beliefs: *safe*, because members of the community can regulate whose testimony they rely on, and *efficient*, because there is a division of epistemic labour whereby members of a community can form true beliefs based on the testimony of good informants, without having to acquire and analyse the informant's evidence. Epistemic evaluations serve as a guide for whom to

⁸Other contributions to this program include Schafer (2014, 2019), Greco and Hedden (2017), Hannon (2019), and Schoenfield (2015). Gibbard (1990) provides a neighbouring account of how the function of rational evaluations (focused primarily on practical rationality) promotes coordination of actions and feelings across a community.

treat as an epistemic surrogate; they allow us to identify whom to *trust* on a particular issue, and perhaps more generally as well (Dogramaci 2012, 524). For example, if I say, ‘Amy’s belief that the berries are poisonous is rational’ or ‘Bel knows whether the prime minister is in Sydney,’ I pick out Amy and Bel as trustworthy epistemic surrogates, and so worthy of being deferred to by members of the community—at least about berries and the prime minister’s whereabouts, respectively. If I were to say that Amy’s belief about the berries is irrational or that Bel doesn’t know the prime minister’s whereabouts, I’d tag Amy and Bel as unworthy of being deferred to on these issues. I’ll call this social role of epistemic evaluation **Deference**.⁹

Deference is the centrepiece of Dogramaci’s account of the social role of epistemic evaluation. But he also thinks it plays other roles in the service of promoting interpersonal coordination. One such role is encouraging norm compliance. He writes that epistemic evaluation carries ‘a brutish kind of force’ (Dogramaci 2012, 522) that encourages others to improve their performance. When I assert that Amy’s belief about the berries is irrational or that Bel doesn’t know the prime minister’s whereabouts, I’m not only marking them as unworthy of being deferred to on these matters, but I’m also *pressuring* them to comply with the norms accepted by the community (Dogramaci 2012, 522). This promotes coordination: as members of a community evaluate one another, performance improves, and more epistemic surrogates are developed. I’ll call this role of epistemic evaluation **Compliance Pressure**.

Here’s the key point for my argument. On Dogramaci’s view, ordinary epistemic evaluations have at least two social roles, both of which aim at modifying the epistemic behaviour of others to promote coordination across a community.

⁹Here’s a potential problem for Dogramaci-style views. We have lots of epistemological evaluative terms—justification, rationality, knowledge attributions, and so on. Sometimes they apply in conflicting ways. Take the brain in a vat who believes they have hands. Most epistemologists accept that BIV rationally believes it has hands yet fails to know it. Justification is more vexed: internalists tend to think BIV’s belief is justified, while externalists tend to think it lacks justification. BIV is positively evaluable in one respect, but not another. Should we think BIV is a reliable informant or not? There are two options here. The first is to follow writers like Craig who take it that just *one* evaluative notion plays the **Deference** role: in his case, knowledge attributions. Dogramaci’s view suggests a second option. He takes it that the role of epistemological evaluations is to solve a coordination problem such that a community can maximize desired outputs (e.g., true beliefs), while minimizing costs (e.g., evidence gathering). Terms need a cost–benefit analysis (Dogramaci 2015, 779–780). The purpose of epistemic evaluation is not just to safely promote true beliefs, but also to do so *efficiently*; their effectiveness needs to be proportional, in some sense, to their efficiency. For instance, some possible evaluative notion according to which one’s belief is positively evaluable iff it’s the result of one’s direct experience is effective but not efficient. I take it that figuring out which notions strike this balance best in a linguistic community will be an empirical question.

Evaluation plays a **Deference** role, regulating how much trust we should place in someone whom we might treat as an epistemic surrogate. And it plays a **Compliance Pressure** role, encouraging underperforming members of a community to do better so that they might become worthy surrogates. Although I've stuck to the details of Dogramaci's account, similar proposals can be found in a number of other accounts of the social role of epistemic evaluation (e.g. Craig 1990, Schafer 2014, Hannon 2019).

What's the upshot for epistemic blame? If Dogramaci-style accounts of the social role of epistemic evaluation are right, then much of the work supposedly done by epistemic blame is already done by epistemic evaluation. **Deference** gives epistemic evaluation the same role that the relationship modification account assigns to epistemic blame. And **Compliance Pressure** gives epistemic evaluation a role that some advocates of the relationship modification account also assign to epistemic blame—for example, Boulton (2021b, 5) claims that epistemic blame aims at 'promoting epistemic goods, like believing truly and avoiding believing falsely.'¹⁰ To the extent that epistemic blame includes commitments to roles like **Deference** and **Compliance Pressure**, it fails to be all that different to epistemic evaluation. This is especially problematic for the relationship modification account, where **Deference** is of central importance. This duplication of roles violates general parsimony considerations, which hold that we shouldn't introduce an idle mechanism that does the same work as an existing one. More specifically, it shows that the relationship modification account fails the test outlined in [Section 1](#): it describes a type of reaction that does not rise above the level of mere evaluation.

How might advocates of the relationship modification account respond? One option might be to claim that the functional equivalence I've identified is merely a verbal dispute: what they call 'blame' Dogramaci calls 'evaluation' and nothing substantive turns on the choice of terminology. If this is right, then Dogramaci's view is nothing for advocates of the relationship modification account to be concerned about since it can be understood as an account of epistemic blame after all, one which might even provide an additional reason

¹⁰Others—who don't necessarily accept the relationship modification account—also assign this role to epistemic blame. Deborah Tollefsen (2017, 362) argues that 'epistemic reactive attitudes' provide a quick and easy way to 'influence conformity' and 'provide motivation to cooperate' in an epistemic community. Likewise, Adam Piovarchy's (2020) view that epistemic blame provides a mechanism for cultivating epistemic agency—such that an agent becomes more sensitive to the expectations of one's epistemic community—assigns epistemic blame a **Compliance Pressure** role. I suspect that everyone who accepts epistemic blame would accept that one of its central roles is discouraging epistemic norm violations.

to be satisfied with their account. I think this objection fails. In the next section I argue that the relationship modification account and Dogramaci's view give different verdicts about a particular class of cases. Not only does this show that the distinction between epistemic blame and evaluation doesn't collapse into a merely terminological matter, but the way in which the relationship modification account deals with these cases provides an additional reason to be unsatisfied with it.

3.2 Deference and excused failings

So far, I've argued that in view of a prominent account of the function of epistemic evaluation, epistemic blame is idle. This violates parsimony and the internal logic of blame itself, which requires blame to rise above mere evaluation. In this section I set out an additional reason to reject the relationship modification account: it is extensionally inadequate as an account of the dynamics of trust. That is, in many cases it gives the wrong recommendations about how to modify trust in response to an epistemic failing. Consider cases of *excused* epistemic failings. In general, blame is only appropriate if an agent culpably violates a norm. If they're excused, we don't blame, but negative evaluation is still appropriate. In these cases, it seems to me that we still ought to lower our trust in the agent even though we don't blame them. Although this result is independent of whether Dogramaci's view of epistemic evaluation is correct, I also think that his view provides more plausible results in these cases. Let's look at an example of excused epistemic failings.

Consider cases of **Diminished Agency**. Suppose I start off with a default, sensible level of trust in several agents in my community. I then learn that they all believe a lot of nutty things and I negatively evaluate them. However, they all have excuses: one is four years old, one was raised in a cult, one had their morning coffee drugged, one has a serious cognitive impairment, and so on. It still seems to me that I should lower my trust in each of these agents as potential informants, even though it wouldn't be appropriate to blame them for the attitudes in question. Perhaps you think that it would be unfair or disrespectful to lower one's trust in someone whose track record of excused epistemic failings is due to diminished agency. But that introduces a new parameter, one which doesn't feature in the dynamics of epistemic trust described by the relationship modification account.¹¹

¹¹You might think that it's just as inappropriate to negatively evaluate such agents as it is to

The relationship modification account holds that if we reduce our trust in someone in response to an epistemic failing, we thereby count as blaming them. But I take it **Diminished Agency** shows that sometimes we reduce our trust in agents in response to their epistemic failings but without blaming them, since we take them to have an excuse. Moreover, **Diminished Agency** highlights that agents can be unworthy informants even if they're blameless. The upshot is that the relationship modification account provides an implausible picture of the dynamics of epistemic trust. Notice that this class of cases isn't a problem for Dogramaci's view; it recommends we lower trust in cases of **Diminished Agency**.

4 The force of epistemic blame

In this section, I focus on the other accounts of epistemic blame: the emotion-based account and Jessica Brown's belief-desire account. I argue that these views fail part of my test outlined in [Section 1](#): they do not provide a plausible account of how epistemic blame carries a 'depth, force or sting,' which is characteristic of blame in general. The views include features that are supposed to satisfy this desideratum, so my argument is not that the force of blame is left out altogether; rather, it's that they give an inaccurate description of our normative epistemological practices. We don't typically respond to others' epistemic failings with the kind of reactions described by these views.

Let's start with the emotion-based account. On this view, negative emotional responses are an essential feature of epistemic blame: we feel resentment, anger, or indignation at others' epistemic failings, and we feel shame, guilt, or embarrassment at our own. The severity of the emotional response supplies the purported force of epistemic blame. I agree with advocates of other theories of epistemic blame who argue that the emotion-based account is at odds with the 'coolness' of our actual practices of responding to epistemic failings (Boult 2020, Piovarchy 2020, Brown 2020b). Usually, we respond to others' epistemic failings without feeling much of anything. The emotion-based

blame them. For example, there does seem to be something strange about telling a four-year-old, or a person living with a serious cognitive impairment, that their belief is irrational. But even if there are constraints on how and when and to whom we express such evaluations, there is still a clear distinction between the appropriateness of blame and evaluation in these cases: it's never appropriate to blame these agents for their epistemic failing, but it's often appropriate to judge and assert that, for instance, their belief is not supported by the evidence or doesn't amount to knowledge.

account predicts that when we learn about an agent's epistemic norm violation—say, they have probabilistically incoherent credences, or believe something on insufficient evidence, or form a belief via an invalid inference—we're disposed to feel frustrated, angry, upset, or feel other negatively valenced emotions towards them.

Speaking for myself, I don't typically have these kinds of emotional reactions to others' epistemic failings, and it doesn't seem plausible to me that others typically do either. This is not to claim that emotions are totally out of place in the epistemological domain. No doubt we sometimes do have emotional reactions to others' epistemic failings. But this is the exception rather than the norm; and in many cases where it seems like emotional reactions to others' attitudes are appropriate, there are other factors at play apart from epistemic norm violation that can explain our reaction. For instance, sometimes the content of a belief *matters* to us in some special way; for example, a Beatles fan might feel angry when she meets people who believe that Paul McCartney died in 1966 and was secretly replaced with a doppelganger. Or it might be that we're connected with the agent in some special (or fraught!) way, or that the attitude makes it more likely that the agent will behave immorally or irrationally, or that we think the agent has a professional obligation to conduct their epistemic life more carefully (say, if they enjoy some political or intellectual influence). Typically, our reactions to purely epistemic failings are not strong enough that the emotion itself carries a depth, force or sting.

Jessica Brown (2017; 2020b, 12–14) offers a diagnosis of why strong negative emotions (and associated dispositions to punish, rebuke, etc) are commonplace in the moral case, but seem out of place in response to epistemological failings. Moral norm violation constitutes a wrongdoing, which usually involves harming a victim. In the epistemological case, it's not clear that norm violation constitutes a wrongdoing, and it usually doesn't involve a victim. Her view is quite agile in the face of these striking disanalogies between moral and epistemological failings. She accepts a variable account of blame where blame needn't include any strong felt emotions and needn't be expressed in practices familiar from the moral domain such as sanctioning. On her view, epistemic blame consists in a characteristic set of dispositions that are causally connected with a belief-desire pair. The belief is that an agent violated an epistemic norm, the desire is that they hadn't, and the dispositions these give rise to—which vary across contexts—are to engage in typical blame-like reactions such as rebuking, asking for justifications, or feeling a negative emotion.

Brown advertises her view's agility as an attractive benefit. Her view retains the core structure of moral blame but adjusts some of its elements in view of the 'coolness' of epistemological reactions. For instance, the *affective strength* of the desire relatum of the belief-desire pair is different in the two domains. In cases of moral blame, people often strongly desire that the target of their blame hadn't acted badly. But in cases of epistemic blame, although Brown takes it that one desires that the agent hadn't believed badly, she holds this desire will usually be much weaker. Likewise for the *dispositions* to feel and behave in characteristic ways that the belief-desire pair gives rise to.

But I think this is a bug, not a feature. Brown hopes to avoid objections about the 'coolness' of epistemic reactions faced by the emotion-based account, but she relies on the same kind of mechanisms as that view to supply the force of epistemic blame. On her view, the force of epistemic blame comes from the desire relatum of the belief-desire pair—the desire that the agent hadn't believed badly—and associated dispositions to feel and behave in certain ways. I think her account faces three challenges. First, it attempts to avoid one problem at the expense of inheriting another. Second, it extends to *other* normative domains in implausible ways. And third, although she holds that both the desire and dispositions to feel in certain ways can be weak, this still seems descriptively inaccurate, so it doesn't avoid the problem it aimed to. Let's look at these challenges.

First, to avoid an account of epistemic blame that has a striking disanalogy with moral blame in terms of how strongly one desires that an agent hadn't violated a norm, she makes a move that lands her with another striking disanalogy between the two species of blame. There's a family of views about moral blame that take it that blame is a type of communication; it signals to others something about the moral norm violation. On many of these views, expressions of blame contain information about the victim of the wrongdoing. Blame might signal that the victim did not deserve to suffer the wrongdoing (Smith 2013, 43), that the victim deserves special care and attention considering the wrongdoing they've suffered (Tierney 2019), that the moral community has good will towards the victim (Kogelmann and Wallace 2018, 9), or that the wrongdoer ought to come to appreciate the severity of their wrongdoing and consider how they might make it up to the victim (McGeer 2012). By conceding that epistemic failings don't constitute a wrongdoing and don't create victims, Brown can explain why we typically don't have strong desires that people hadn't believed badly, but she inherits a new disanalogy that many will

find just as striking a contrast between moral and epistemic blame.¹²

Perhaps Brown might reply that this criticism depends on features of moral blame that she rejects; after all, she's working with Sher's account rather than communicative accounts. In that case, this objection might only be a conditional worry: if communicative theories of blame are apt, then Brown's view faces the above challenge. However, epistemologists who accept the notion of epistemic blame often motivate it by appealing to communicative elements, so I think this challenge should be addressed directly. For instance, the intuitive case for epistemic blame is *strongest* when it seems like violating an epistemic norm has a detrimental effect on others or their informational environments. Our reactions are quite different depending on whether it seems like an epistemic failing might cause real damage. My reaction to a student who makes several epistemic failings in their paper—say, an insensitivity to the evidence and an inability to adequately justify their view—is quite different to a politician who makes the same failings when defending delaying action on climate change or spreading misinformation about election fraud. In the latter case, it seems much more plausible that my reaction *might* be a species of blame. So even though Brown opts for an account of blame that doesn't involve communicating anything about a wrongdoing or a victim, some of the most promising examples used to motivate the very idea of epistemic blame rely on these elements.

Second, Brown claims that Sher's account of moral blame can be adapted for the epistemological domain. If this is correct, it's natural to think this is because it's getting something right about blame in general such that it could provide a plausible account of blame in other normative domains too. But if we try to extend the account in this way, I think we get some implausible results. Take practical rationality. Sometimes we have negative reactions to others' prudential decisions, but it's unnatural to think of these as a species of *blame*. But when Brown's account is extended to practical rationality these reactions are a kind of blame after all. In fact, her account generates distinctive kinds of blame for all domains that include norms governing actions and attitudes; all that's required is a belief-desire pair that's casually connected with dispositions to feel and engage in blame-like responses—and recall that the desire and associated dispositions can be weak. So, Brown's account faces a challenge: explain why we ought to stop with just moral and epistemic blame, or defend her ac-

¹²I'm grateful to Hannah Tierney for discussion on this point.

count's results about distinctive kinds of prudential blame, aesthetic blame, and so on.

Here's a case that illustrates this objection. Suppose my friend regularly makes poor life choices: he lives a very unhealthy lifestyle, never saves for retirement, takes on more and more debt, and sabotages himself before he can complete any important project. His life is filled with prudentially tragic choices. Presumably I believe he's violated a norm of practical rationality—such as failing to maximize expected utility (or whatever you think are the correct norms of prudence)—and I desire that he hadn't. Suppose the belief-desire pair cause dispositions to feel disappointed with him or engage in practices like telling him to get his act together or lowering my estimation of him. All this fits the bill of blame on Brown's view. But it just seems implausible to me that this amounts to blaming my friend in a distinctly prudential sense. Of course, Brown never claims that her account of blame is fully general. But the fact that her account is implausible when extended to other domains gives us reason to doubt its claims about epistemic blame.¹³

Third, I think Brown's view fails the same test as the emotion-based account. She claims that epistemic blame involves desiring that an agent hadn't believed badly. Although she claims that this desire can be weak (Brown 2020b, 400–401)—and the associated feelings and behaviours it gives rise to can be weak and varied—it's still going to have to be strong enough that the desire *itself* supplies some force or that the feelings/behaviour it disposes one to carry force. Weak desires can accompany mere evaluation, where the latter doesn't constitute blame. Suppose I criticise Caligula as immoral and weakly desire that he'd lived a better life. It doesn't seem that adding in the weak desire changes my criticism from mere evaluation to blame. If the desire relatum of the belief-desire pair is going to supply epistemic blame with some force, it's going to have to have a certain level of affective strength. Desires of this kind are as descriptively inaccurate of our practices of epistemological criticism as the strong negative feelings posited by the emotion-based account.

Things get murky if we try to pin down a threshold for what counts as a

¹³Even if one is sympathetic to the notion of blame in other domains, Brown still faces the challenge of defending the particular accounts of, say, aesthetic or prudential blame that her view generates. These accounts are highly revisionary of our ordinary normative practices within these domains, classifying a wide range of mundane reactions as blame. So I take it those sympathetic to Brown's view cannot avoid this objection just by accepting the general idea of blame in various domains; they need to show that Sher's framework can be used to develop plausible accounts of various species of blame. I'm grateful to an anonymous reviewer for helpful comments on this point.

non-trivial desire. Nevertheless, I think it's plausible that desires of roughly this strength are rarely a feature of our typical reactions to others' epistemic failings. Take the example of inconsistent beliefs David Lewis (1982, 436) mentions in his discussion of fragmentation. He says he used to believe that Nassau Street ran roughly east-west; that the railroad nearby ran roughly north-south; and that the two ran roughly parallel. Do I desire that he hadn't believed badly? No. Am I disposed to feel much of anything about his inconsistent beliefs? No. And even if I am on some level that escapes my conscious awareness, it's only in a very weak sense. As with the emotion-based account, I don't claim that we never have stronger desires or dispositions to feel certain ways in response to others' attitudes. My answer to each question changes if we modify the case in certain ways—for instance, if Lewis had confessed to having contradictory beliefs about some of his influential views, or about issues related to social policy. But as with the emotion-based account, this introduces other elements, which suggests that when we do have a strong desire that an agent hadn't believed badly, or are disposed to feel or react in ways that carry some opprobrium, it's not due to a bare epistemic failing.¹⁴

Perhaps intuitions diverge here. But at this point, the breadth of the blameless/blameworthy distinction in epistemology counts in my favour. Writers who accept epistemic blame take it that it captures how we *actually* respond to others' epistemic failings across a wide range of cases. If there's purely epistemic blame as Brown describes it, a non-trivial level of desire that an agent hadn't believed badly—which gives rise to dispositions to have non-trivial feelings or other kinds of reactions—is a systematic feature of our reactions to others' epistemic failings. The fact that we rarely react to people in this way counts against the notion of epistemic blame. Perhaps Brown might reply that although we don't typically blame others for their epistemic failings in these ways, that we *ought* to (or at least are permitted to). But this makes her view quite revisionary of our actual practices. This leaves her with the challenge of justifying why we ought to turn the temperature up on these reactions, and it's not obvious what benefits this would have for either our epistemic or social lives. Overall, Brown's view faces the same challenge as the emotion-based

¹⁴It bears emphasizing that there is a striking asymmetry with moral blame here. In the epistemological case, our reactions to others' epistemic failings rarely include a non-trivial level of desire that an agent hadn't believed badly; and when they do, that's because non-epistemological factors are at work. The same is not true in the moral case where our reactions to others' moral failings often involve a non-trivial desire that the agent had behaved differently. I'm grateful to an anonymous reviewer for helpful comments on this point.

view: both views account for the force of epistemic blame by making implausible claims about the psychology of epistemological criticism.

5 Conclusion

I've advocated scepticism about epistemic blame. I've defended this view by arguing that all extant accounts of epistemic blame fail to pass a simple test for what it takes for a kind of reaction to count as a species of blame. My view has two upshots. First, it uncovers an interesting fault line between epistemic and moral normativity: we have good reason to think there's no purely epistemic kind of blame. Second—and more importantly—it rules out a common strategy for defending highly demanding epistemic norms. If there's no purely epistemic kind of blame, classifying epistemic failings as either blameworthy or blameless is inapt.

Fans of demanding epistemic norms might respond that they're happy to give up the notion of epistemic blame since there are other ways to carry out the same kind of strategy. For instance, perhaps notions of excuse or culpability will serve just as well for their purposes of responding to demandingness objections. But in other normative domains, these categories are so tightly connected with notions like blame, sanction, and punishment, that it's not clear that drawing a distinction between, say, an excused and culpable epistemic failing has much normative significance. For example, if there's no epistemic blame, it's puzzling what one is excused *from*. Giving up the notion of epistemic blame has an impact on the wider normative landscape that prevents the strategy being pursued simply by shifting to a neighbouring notion.

References

- Basu, Rima. 2018. "Can Beliefs Wrong?" *Philosophical Topics* vol. 46, no. 1: 1–17.
- .2019a. "What We Epistemically Owe To Each Other." *Philosophical Studies* vol. 176, no. 4: 915–931.
- Basu, Rima and Mark Schroeder. 2019. "Doxastic Wrongings." In *Pragmatic Encroachment in Epistemology*, eds. Brian Kim and Matthew McGrath. Routledge, pp. 181–205.

- Boult, Cameron. 2020. "There is a Distinctively Epistemic Kind of Blame." *Philosophy and Phenomenological Research*. Early view, doi: 10.1111/phpr.12726
- .2021a. "Standing to Epistemically Blame." *Synthese* Online First.
- .2021b. "The Significance of Epistemic Blame." *Erkenntnis* Online First.
- Brown, Jessica. 2017. "Blame and Wrongdoing." *Episteme* vol. 14, no. 3: 275–296.
- .2020a. "Epistemically Blameworthy Belief." *Philosophical Studies* vol. 177: 3595–3614.
- .2020b. "What is Epistemic Blame?" *Noûs* vol. 54, no. 2: 389–407.
- Christensen, David. 2004. *Putting Logic in Its Place*. Oxford University Press.
- Craig, Edward. 1990. *Knowledge and the State of Nature*. Oxford University Press.
- Dogramaci, Sinan. 2012. "Reverse Engineering Epistemic Evaluations." *Philosophy and Phenomenological Research* vol. 84, no. 3:513–530.
- .2015. "Communist Conventions for Deductive Reasoning." *Noûs* vol. 49, no. 4: 776–799.
- Gendler, Tamar Szabo . 2011. "On the Epistemic Costs of Implicit Bias." *Philosophical Studies* vol. 156, no. 1: 33–63.
- Gibbard, Allan. 1990. *Wise Choices, Apt Feelings*. Harvard University Press.
- Greco, Daniel. Forthcoming. "On the Very Idea of an Epistemic Dilemma." In *Essays on Epistemic Dilemmas*, ed. Nick Hughes. Oxford University Press.
- Greco, Daniel and Brian Hedden. 2017. "Uniqueness and Metaepistemology." *The Journal of Philosophy* vol. 113, no. 8: 365–395.
- Hannon, Michael. 2019. *What's the Point of Knowledge? A Function-First Epistemology*. Oxford University Press.
- Hawthorne, John and Amia Srinivasan. 2013. "Disagreement Without Transparency: Some Bleak Thoughts." In *The Epistemology of Disagreement: New Essays*, eds. Jennifer Lackey and David Christensen. Oxford University Press, pp. 9–30.
- Hawthorne, John and Jason Stanley. 2008. "Knowledge and Action." *The Jour-*

- nal of Philosophy* vol. 105, no. 10: 571–590.
- Hieronymi, Pamela. 2004. “The Force and Fairness of Blame.” *Philosophical Perspectives* vol. 18, no. 1: 115–148.
- Kogelmann, Brian and Robert H. Wallace. 2018. “Moral Diversity and Moral Responsibility.” *Journal of the American Philosophical Association* vol. 4, no. 3: 371–389.
- Lasonen-Aarnio, Maria. 2010. “Unreasonable Knowledge.” *Philosophical Perspectives* vol. 24, no. 1: 1–21.
- Lewis, David K. 1982. “Logic for Equivocators.” *Noûs* vol. 16, no. 3: 431–441.
- McGeer, Victoria. 2013. “Civilizing Blame.” In *Blame: Its Nature and Norms*, eds. D. Justin Coates and Neal A. Tognazzini. Oxford University Press, pp. 162–188.
- McHugh, Conor. 2012. “Epistemic Deontology and Voluntariness.” *Erkenntnis*, vol. 77, no. 1: 65–94.
- Nottelmann, Nikolaj. 2007. *Blameworthy Belief: A Study in Epistemic Deontology*. Springer.
- Peels, Rik. 2013. “Against Doxastic Compatibilism.” *Philosophy and Phenomenological Research* vol. 89, no. (3): 679–702.
- Piovarchy, Adam. 2020. “What do We Want from a Theory of Epistemic Blame?” *Australasian Journal of Philosophy* vol. 99, no. 4: 791–805.
- Reid, Thomas. 1785/2002. *Essays on the Intellectual Powers of Man*. Penn State University Press.
- Rettler, Lindsay. 2018. “In Defence of Doxastic Blame.” *Synthese* vol. 195, no. 5: 2205– 2226.
- Scanlon, T.M. 2008. *Moral Dimensions: Permissibility, Meaning, Blame*. Belknap Press.
- .2013. “Interpreting Blame.” In *Blame: Its Nature and Norms*, eds. D. Justin Coates and Neal A. Tognazzini. Oxford University Press, pp. 84–100.
- Schafer, Karl. 2014. “Doxastic Planning and Epistemic Internalism.” *Synthese* vol. 191: 2571–2591.
- .2019. “Epistemic Planning, Epistemic Internalism, and Luminosity.” In

- Metaepistemology*, eds. Conor McHugh, Jonathan Way, and Daniel Whiting. Oxford University Press, pp. 179–195.
- Schoenfield, Miriam. 2015. “Internalism Without Luminosity.” *Philosophical Issues* vol. 25, no. 1: 252–272.
- Schmidt, Sebastian. 2021. “Epistemic Blame and the Normativity of Evidence.” *Erkenntnis* online first.
- .2022. “Blameworthiness for Non-Culpable Attitudes.” *Australasian Journal of Philosophy*. DOI: 10.1080/00048402.2022.2055093
- Sher, George. 2006. *In Praise of Blame*. Oxford University Press.
- Smith, Angela. 2013. “Moral Blame and Moral Protest.” In *Blame: Its Nature and Norms*, eds. D. Justin Coates and Neal A. Tognazzini. Oxford University Press, pp. 27–48.
- Tierney, Hannah. 2019. “Expanding Moral Understanding.” *Australasian Philosophical Review* vol. 3, no. 3: 318–323.
- Tollefsen, Deborah Perron. 2017. “Epistemic Reactive Attitudes.” *American Philosophical Quarterly* vol. 54, no. 4: 353–366.
- Watson, Gary. 2004. *Agency and Answerability*. Clarendon Press.
- Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford University Press.
- .2005. “Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge.” *Philosophical Quarterly* vol. 55, no. 219: 213–235.
- .2013. “Response to Cohen, Comesaña, Goodman, Nagel and Weatherson on Gettier Cases in Epistemic Logic.” *Inquiry* vol. 56: 77–96.
- .2014. “Knowledge First.” In *Contemporary Debates in Epistemology*, eds. Matthias Steup, John Turri, and Ernest Sosa. Wiley-Blackwell, pp. 1–9.
- .Forthcoming. “Justification, Excuses, and Sceptical Scenarios.” In *The New Evil Demon*, eds. Julian Dutant and Fabian Dorsch. Oxford University Press.