




ARTICLE

Scepticism About Epistemic Blame Scepticism

Adam Piovarchy 

The University of Notre Dame, Australia

Email: adam.piovarchy@nd.edu.au

(Received 7 August 2024; revised 22 December 2024; accepted 1 January 2025)

Abstract

A number of philosophers have recently argued that there is such a thing as ‘epistemic blame’: blame targeted at epistemic norm violations qua epistemic norm violations. However, Smartt (2024) and Matheson and Milam (2022) have recently provided several arguments in favour of thinking epistemic blame either doesn’t exist or is never justified. This paper argues that these challenges are unsuccessful and along the way evaluates the prospects for various accounts of epistemic blame. It also reflects on the dialectic between sceptics and realists about epistemic blame and what choice points are available for moving the debate forward.

Keywords: Blame; responsibility; epistemic blame; excuses; doxastic responsibility

1. Introduction

A number of epistemologists have recently argued that there is such a thing as ‘epistemic blame’: blame targeted at epistemic norm violations qua epistemic norm violations (Brown, 2020a, 2020b; Boulton, 2021a, 2021b, 2023, 2024; Piovarchy, 2021; [forthcoming-a](#); Rettler, 2018). Although this attention could be interpreted as an exciting new discovery of a hitherto unanalysed phenomenon, others are much more sceptical. Matheson and Milam (2022) think that all non-moral forms of blame can never be justified, while Smartt (2024) goes further, arguing that in fact ‘there is no distinctively epistemic form of blame’ (p. 2).

These philosophers’ challenges have multiple fronts, raising a number of issues that have not yet received attention in the epistemic blame literature. I, however, believe that epistemic blame is up to the task. It is worth considering how epistemic blame realists might respond to such sceptical challenges because our stance on this topic has wide-ranging implications. Epistemic blame has notably been used to defend externalist approaches to justification against charges of demandingness, by allowing proponents to argue that while subjects lack justification, they are blameless (Williamson, [forthcoming](#), cf. Brown, 2020; Ballarini, 2022). Epistemic blame has also been thought relevant to thinking about epistemic normativity more broadly (Boulton, [forthcoming](#)), the normativity of evidence (Schmidt, 2021), and attempts at identifying distinctly epistemic norms (Kaupinnen, 2018).¹ One way to respond is by shoring up existing

¹Though Kaupinnen talks of holding people epistemically accountable, rather than blaming.

positive proposals, arguing that sceptics cannot undermine some particular theory. Another is to undermine the reasons for doubting epistemic blame exists. This paper primarily takes the latter approach.² I believe Smartt and Matheson and Milam's arguments face some challenges themselves, and we should consider replies available to epistemic blame realists. Doing so will also serve as an opportunity for evaluating the prospects of different theories of epistemic blame. Since there are now quite a few live accounts of epistemic blame, it is worth considering their strengths, weaknesses, and seeing if any are particularly viable. This will also act as an opportunity to clarify how theorising about epistemic blame connects up with our actual practices, and what does or does not need to be accounted for when intervening in this debate.

This paper proceeds as follows. First, I consider Matheson and Milam's (2022) arguments that non-moral blame is never justified, because we do not have epistemic obligations. They attempt to show this by arguing that agents' blame in key thought experiments is always moral, or unfitting. I show that their approach at best amounts to reporting certain intuitions and denying the existence of others' intuitions, a strategy which is unhelpful for adjudicating the debate. I then consider Smartt's argument that Brown's (2017, 2020) treatment of epistemic blame misrepresents our interactions with people who violate epistemic norms, arguing that the epistemic blame proponent can explain all putative discrepancies. I then consider Smartt's argument that evaluations can do the work that epistemic blame purportedly does and argue that they do not. I finally consider how epistemic blame realists can prevent other kinds of blame from being admitted into their ontology, noting that an account I've (2021) previously defended is unscathed. I conclude by summarising what appears to be the dialectic between epistemic blame realists and sceptics and where sceptics need to focus their efforts if they wish to gain the upper hand.

2. Moral and non-moral targets

Matheson and Milam (2022) attempt to show that non-moral blame is never justified, in that it is never a morally permissible response to a non-moral failure. They agree with the realist that people do exhibit blaming reactions in response to seemingly non-moral failures and note that our existing practices of non-moral blaming have several similarities to our moral blaming practices. Given this, the challenge is to identify a disanalogy between the two. Their approach is to show that, since there is 'a conceptual connection between blameworthiness and obligation such that one is blameworthy only if one fails to meet an obligation' (204), there are no non-moral obligations. Their method for achieving this in turn is to argue that all cases of putative non-moral blame are actually either moral blame, or unfitting.³

Their first line of argument for thinking that many instances of non-moral blame are, in fact, moral is to note that 'when challenged about the permissibility of their non-moral blame, blamers often fall back on different reasoning' (p. 205). For example, when

²Though §5 will also involve the former.

³More specifically, they initially argue that non-moral blame is *often* either moral or unfitting, and only later then move to challenging the realist to find a case in which the target of blame exhibits a genuinely non-moral failure while meeting various other conditions, and also explain the obligation that was violated. This (conveniently) doesn't show non-moral blame doesn't exist, rather, it functions more as an attempt to shift the burden-of-proof onto the epistemic blame realist to show non-moral blame exists according to the sceptic's standards, without assessing whether these standards are reasonable terms for the debate. For simplicity and to highlight why this attempted burden-shifting is unsuccessful I'll just directly examine their handling of various cases and argue this fails to shift the burden.

asked why a missed goal warrants athletic blame, a blamer might point to things such as the player's huge salary, the fact that he should have practiced more or that his shot demonstrates a lack of care, all of which Matheson and Milam think are moral grounds.

The efficacy of this line of reasoning is dubious because most people are not adept at distinguishing questions of the circumstances that are sufficient for blame to be fitting with questions of what justifies blaming in those circumstances (indeed, Matheson and Milam do not make this distinction either). In response to 'Why are you blaming him?', the reply 'Because he missed the shot!' is a perfectly intelligible and fitting answer. What justifies blaming poor performances more generally is another question altogether, and most people are not particularly well-versed in the range of different options one might consider when trying to ground different types of obligations or blaming responses.⁴ What would be needed is for the sceptic to show that no such grounds are available, and that would require engaging with questions of the foundations of normativity. (We shall return to this question in §5.) Consequently, this line of argument does not succeed against the epistemic blame realist. That a moral consideration in favour of blaming *can* be identified as present doesn't thereby show that *that* consideration is what is doing all the work.

Teasing out what is doing the work is attempted in Matheson and Milam's second line of argument, which is that 'our attitudes seem to change when we strip away the moral features of a particular case' (p. 205). If we consider cases in which someone makes a poor artwork, or engages in shoddy reasoning, but in which nothing significant turns on this particular performance, Matheson and Milam think that it seems harder to sustain an intuition that the agent is blameworthy.

However, this method threatens to bottom out in a mere trading of intuitions, which we can see by imagining arguing with someone who was convinced that *only* non-moral blame was ever justified. (Perhaps they had been inspired by Socrates' argument that no-one ever knowingly does evil.) Suppose you encountered a community of philosophers who were convinced that all justifiable blame only ever targeted epistemic, aesthetic, or athletic (etc.) failures. There are plenty of such considerations to go around, e.g. the rationality of committing a murder often involves hasty reasoning, and people find dead bodies to be aesthetically grotesque. Suppose now you tried to convince them that their worldview was mistaken, and that blame actually targets moral failings, not non-moral ones. They would worry that your assessments were being confounded; that while you took yourself to be blaming moral failings, you were actually responding to these other types of failings. Given this, they would reason that to test your theory, we would need to consider cases involving *purely* moral failings with *no* non-moral considerations present whatsoever.

I suspect there would not be many cases both sides could agree meet this criteria, and of those that remained, it wouldn't be difficult for these philosophers to report an intuition that blame didn't seem especially warranted. Of course, we would find their position ludicrous, and the considerations that they pointed to as 'really doing the work' to be mistaken. But this shows such a line of argument doesn't do much to advance the debate, so much as merely shore up one's pre-existing position. Because the non-moral blame realist takes both non-moral *and* moral blame to be justified, the method of trying to fully 'control' for confounding factors and then restricting our attention to a very

⁴Rawls (1955) famously distinguishes between justifying a practice and justifying an action under that practice. While we might have a system of punishment in order to deter (a forward-looking consideration), this doesn't mean punishment of individuals is only fitting when it will deter – we instead punish because they broke the law (a backward-looking consideration). Use of such a distinction is how I've previously argued (2021) we can find a justification for epistemic blame.

narrow range of cases in which it appears uncostly to report a particular intuition is something that only the sceptic can propose. Showing why this method is suspect requires seeing how it looks when made by someone who rejects a type of grounds for blameworthiness that we accept, which our imagined community of moral-blame-insensitive philosophers helps illuminate.

If we want to make some dialectical progress in establishing what considerations are doing the work, a viable method would be to reflect on our actual practices, or find some independent consideration from which to establish some kind of burden-of-proof. But these both seem to favour the non-moral blame realist: as Matheson and Milam note, people *do* seem to blame others for non-moral failings, and more importantly frequently take themselves to be doing so. The non-moral blame realist thus seems to better capture the phenomenology of our blaming interactions. When we blame an athlete for, say, missing the goal, we are blaming him *for his poor performance*. When the athlete then apologises, they don't apologise for some kind of wrongdoing or moral transgression, they apologise *for missing the goal*. This is enough to mean that the sceptic's interpretation of our practices, then, is somewhat revisionist.⁵

Because appealing to intuitions about the blameworthiness of agents in particular cases (or the blameworthy-making features of such agents) threatens to result in a stalemate between the epistemic blame realist and sceptic, an alternative approach is needed. This is what Smartt (2024) attempt to provide, presenting a number of arguments against epistemic blame realism, which it will be useful to consider individually.

3. Force and frequency

Smartt's first argument is that attempts to explain what force is provided by epistemic blame end up with an implausible picture of our practices of epistemic criticism. Here he focuses on Brown's (2017, 2020) belief-desire account of epistemic blame, itself built upon Sher's (2005) account of moral blame. Brown takes blame to consist in a belief that the agent has believed badly, a desire that they had not believed badly, and a range of behavioural and affective dispositions which this belief-desire pair then manifests. When I blame someone for having inconsistent beliefs without a good excuse, for instance, I believe that they have believed badly, I desire that they had not done this, and this results in a range of dispositions (e.g. disappointment and reduction of trust) which count as blame. Smartt's target, however, is broader than this. In the moral blame literature, a large number of accounts take blame to involve some kind of negative

⁵Matheson and Milam also offer several additional arguments for their conclusion, but they are easily refuted. They note that angry blame seems impermissible when no moral considerations are present, and infer from this that non-moral failures are not blameworthy. But not all blame features anger, and all that is needed for the non-moral blame realist to succeed is to identify one case in which such blaming responses remain fitting. They also attempt to bolster this argument by comparing 'two cases, identical but for the fact that one has a moral dimension' (p. 206), but this method is flawed. That the added moral dimension makes us blame more does not show that the agent without the moral dimension is blameless. (By analogy: that unexcused killing is more blameworthy than unexcused stealing does not mean stealing is not blameworthy). Finally, they argue that we often fail to meet various condition on permissible blaming, namely that we may not be warranted in believing that the target is blameworthy due to limited evidence, we may lack standing to blame, our blame may be unfair, and because non-moral blame is pointless and fails to produce good outcomes. However, these are all also true of many cases of moral blame, without thereby coming close to showing that moral blame is never permissible. For arguments that epistemic blame is not universally hypocritical, see Piovarchy (forthcoming). For arguments that non-moral blame can serve important social goals, just like moral blame, see §5.

emotional reaction towards wrongdoing. Since it is plausible that we experience a range of negative emotional reactions towards agents who violate epistemic norms (*qua* epistemic norms) without an excuse, it is plausible that this is a distinct type of blame. My own (2021) account of epistemic blame similarly takes certain emotional responses (e.g. disappointment, frustration, non-retributive anger) to be forms of epistemic blame, and so any objections to emotion-based accounts will undermine this account too.

Smartt's objection concerns how well claims about our epistemic blaming responses map onto our interactions with putatively epistemically blameworthy agents. He argues 'The emotion-based 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' (p. 1816). But, he notes, most epistemic norm violations don't generate these kinds of emotions. And in those cases where negative emotions do seem appropriate, factors beyond epistemic blame can explain our reactions.

Epistemic blame realists can respond to this challenge, but let's first distinguish two questions that seem to be being run together: why epistemic blame seems, on the whole, less *frequent* than moral blame, and why certain instances of epistemic blame seem less *fitting* than moral blame.⁶ Accounting for the reduced frequency of epistemic blame is simple for realists: in those circumstances where we don't blame, we didn't care about those particular violations (or at least, didn't care enough to experience any blame). These accounts of epistemic blame don't *predict* anything about how frequently we will experience blame upon learning of a norm violation, and nor do they need to. The same obtains for moral norm violations: I might fail to experience moral blame in response to hearing about the latest political scandal on the news, but that doesn't mean that the perpetrators aren't morally blameworthy, or even that I don't judge them morally blameworthy. It's just that, for whatever reason, I don't particularly care about them; perhaps I am disenchanted with politics. If it ever came to arguing that moral blame exists, what would matter is that there are plenty of cases where I do, in fact, experience moral blame, and that I have a principled reason for thinking that people can be morally blameworthy in those cases. A similar move can thus be made concerning the frequency of epistemic blame.

When we turn to examining the blame that people *do* feel (and its fittingness), Smartt thinks there are problematic disanalogies between moral and epistemic blame, in that only the former involve a victim and strong negative emotions. But it seems to me that if anything, the analogy holds: given there's a difference in whether there's a victim, we should expect some difference in the content and affective strength of moral and epistemic blame, and this is precisely what we observe. There's a *difference* in how strongly we experience the reactions, but it isn't *prima facie* problematic as there are also differences between kinds of moral norms and how we respond to violations of those. Some moral norm violations are so trivial that they also don't generate much resentment, but this isn't a problem for thinking of the emotional reactions that do occur as a form of moral blame.

When we try to explain away the difference in strength between moral and epistemic blame by pointing to the lack of a victim, Smartt worries that only we end up with another disanalogy, namely that moral norms communicate or signal information about

⁶E.g. 'Brown 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'. [emphasis added]. 'Commonplace' suggests concern for frequency, but 'out of place' seems to target fittingness. Claims about weakness and strength would target fittingness.

victims, whereas epistemic norms don't, despite epistemic blame realists sometimes motivating their account by appealing to the communicative elements of blame. But they need not take it to be an essential feature of blame that it involves concern for a victim. Why not think that blame targets and communicates disapproval of norm violations, and that moral norm violations often involve victims while epistemic norms don't? Identifying that there is a disanalogy here between moral and epistemic blame doesn't point to any problem, unless we should independently think that blame needs to communicate information about victims, or unless we should think that blame doesn't target norm violations.⁷

We can weaken Smartt's framing by noting that many moral norm violations simply have no victim, and yet still warrant moral blame. Perhaps the wrongdoing was an act that conflicted with some abstract value that generates substantive duties, or violated something like the categorical imperative. Perhaps the agent failed to pull their weight in a collective action problem, but success with the problem was over-determined. Perhaps the agent had no effect on the victim, but acted with poor quality of will, such as sadistic voyeur. Perhaps in acting wrongly, the 'victim' was actually benefitted and was grateful that the agent acted wrongly!

In response to worries that we don't need our account of blame to involve communicating concern for victims, Smartt notes that the case for epistemic blame seems strongest when the norm violation does have negative effects on other agents or their informational environment. But arguably it only seems strongest insofar as we lack an account of epistemic blame's frequency and fittingness. Once this is on the table, there is little left to explain, and proponents will say the case for epistemic blame realism seems stronger. We already saw we can account for why people don't frequently experience epistemic blame by saying they don't care about those particular violations. We can likewise account for the cases Smartt thinks are strongest by saying they happen to be the instances in which people care more about the violations in question. It is indeed a legitimate question *why* we care about these cases more. Epistemic blame proponents can accept that such cases feature moral blame too, or that some epistemic norm violations are more important than others. Sceptics, in turn, might argue that moral considerations and moral blame are doing all the work, but epistemic blame realists are aware that there are multiple factors here potentially affecting our intuitions and have provided careful analyses of thought experiments to make their case (e.g. Piovarchy, 2021, p. 792–793; Brown, 2020 pp. 2–3, 12–13; Boulton, 2021, p. 521; 2023, p. 819; Goldberg, 2017, §4). As we saw with Matheson and Milam, if sceptics want to do more than report a different set of intuitions, they need to either engage with the realists' treatment of these cases or hang their argument elsewhere.

Smartt's first objection, then, doesn't necessarily stick. However, the epistemic blame sceptic might nevertheless be on the road to their goal, because Brown's account is known to face a number of objections. In particular, it seems one can desire that others not believe badly and have this manifest in certain ways without being blamed, one can blame without having the desire that someone not have believed badly, and there are some components of Sher's account of moral blame that cannot be transplanted to epistemic blame (Boulton, 2023, §2; Piovarchy, 2021, pp. 797–799). If there is a way to also

⁷Putting this in terms of fittingness also won't succeed. Typically, to say that an emotion is fitting is to say that it correctly represents its object. To say that blame is fitting is typically understood as saying that blame correctly represents its target as blameworthy, i.e. as a culpable wrongdoer. If Smartt is proposing that blame instead necessarily represents victims as needing concern, this will require further argumentation. It is worth noting that the communicative accounts that Smartt points to in support of this line of argument do not claim that blame *invariably* or *essentially* communicates concern for victims.

undermine competing accounts of epistemic blame, namely those from Boulton and myself, the sceptic might succeed yet. Fortunately, Smartt has more objections.

4. Co-ordination and redundancy

Smartt's next argument is that much of the work that epistemic blame is supposed to do is already carried out by epistemic evaluation, which plays two important roles. First, evaluation enables *Deference*. In saying 'Bel knows whether the Prime Minister is in Sydney', I mark Bel as a trustworthy epistemic surrogate, and thus worthy of being deferred to by others in the epistemic community. Second, evaluation provides *Compliance Pressure*. When I mark Bel's belief as irrational, I not only mark her as untrustworthy, I pressure her to comply with the norms of the community. This, in turn, enables co-ordination, as their performance improves.

Smartt argues this particularly poses a problem for Boulton's (2021a, 2021b, 2023) account of epistemic blame, which is built off of Scanlon's (2008) account of moral blame. The details of Boulton's account – which have several moving parts – do not actually matter for understanding and engaging with this objection.⁸ All that matters is that Boulton takes epistemic blame to aim at 'promoting epistemic goods, like believing truly and avoiding believing falsely' (2021, p. 5).

Smartt (2024, fn. 10) also has my (2021) account in his sights with this objection. On my account (itself building on Vargas's (2013) account of moral blame), epistemic blame starts off as an external motivator (pressure) to comply with epistemic norms, but gradually, over time, causes agents to internalise those norms. Epistemic blame helps to cultivate a distinctively valuable kind of agency, one that is responsive to epistemic norms. The difference between moral and epistemic blame's natures is accounted for by the fact that the norms we are trying to sensitise other agents to in blaming to are quite different in character, with epistemic norms notably lacking any victims, but both rely on blame's unpleasant features to achieve their goal.

Smartt's complaint is that since epistemic evaluation already enables deference and provides compliance pressure, then epistemic blame is not adding anything. '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' (p. 1819–1820). This particularly means that the accounts of epistemic blame from Boulton and myself fail to provide any required force or sting, commonly taken to be a desideratum on a theory of blame.⁹

⁸In case readers are curious: on Scanlon's account, to blame is to modify the intentions and expectations that constitute our epistemic relationship with another agent. A relationship is constituted by a set of expectations and intentions towards the other person, according to some normative ideal. To judge that someone has done something faulty by the standards of a relationship is to judge them blameworthy. To then modify one's relationship (understood as the set of intentions and expectations) in a way that said judgment would render fitting, is to blame.

Boulton similarly takes there to be a general epistemic relationship between all agents, which has epistemic trust as its normative ideal. To say that you are in a relationship of epistemic trust '... 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... When all goes well—that is to say, according to the normative ideal of this relationship of mutual epistemic trust—each member intends to epistemically trust another unless one has good reason not to'. (Boulton, 2023, p. 816). To modify our intentions and attitudes towards someone on the basis of their impairing our relationship is to blame them.

⁹Smartt may be correct that a relationship modification account of epistemic blame fails to provide the relevant force we want from a theory of blame, but this may have less to do with epistemic blame and more

Smartt's argument, I believe, is too hasty in its reasoning. Showing that evaluations *can*, in principle, play this role to some degree is not the same as showing that they *alone* do, in practice, play that role sufficiently.¹⁰ We can see this by going back to the moral case. Moral evaluations *also* provide compliance pressure. Claiming that someone did the wrong thing, or is evil, or callous, or negligent, exerts some form of social pressure. But it would clearly be too hasty to infer that as a result of this, moral blame does not add anything beyond what is already provided by moral evaluations. Responding to a wrongdoer by simply evaluating them seems like a poor method for making them understand the significance of their wrongdoing and motivating them to apologise. It is one thing to show that epistemic evaluation provides pressure, but another to show that it provides pressure to the degree and in the same manner as existing accounts of epistemic blame.

This reply seems even stronger when made on behalf of the agency-cultivation account, which gives a more explicit role to compliance pressure. To be sure, there are many factors which *can* improve the chances that an agent will internalise norms. But we cannot infer from this that all factors work equally well in all contexts, which is what would need to be shown for charges of redundancy to stick.

Additionally, showing that epistemic blame's role is already played by something else is not the same as showing that epistemic blame does not exist. To get to the latter, Smartt argues that we should prefer epistemic blame scepticism on grounds of parsimony, avoiding introducing idle wheels that do the same work as existing ones. But there is an important ambiguity here in what it means to 'do the same work'. Epistemic evaluation and epistemic blame might 'do the same work' in the sense of both potentially being able to achieve the same end, given our collective interest in designing systems and practices that result in agents complying with epistemic norms. But parsimony is usually invoked in explanations or models which are aiming to account for our observations, and we observe some important differences in how people evaluate and how they blame, which count against thinking that evaluation alone can do all the work. Smartt elides a distinction in how it is that *Deference* produces coordination across an epistemic community. Evaluations primarily seem to be for the benefit of *others*; when I call someone irrational, I am letting others know not to rely on them. But when I engage in epistemic blame, modifying my relationship, the primary change is *my* willingness to rely on them. That these can come apart is clear, because I can sincerely call and evaluate someone as irrational, and yet still subsequently rely on them without experiencing any blame. Alternatively, I can epistemically blame someone, and yet refrain from making any public evaluations of them whatsoever. If epistemic blame is an idle mechanism, then Smartt needs to invoke some other mechanism to explain why these responses appear to be different, or show they are the same response simply described in different ways.

That evaluations of blameworthiness (whether moral or epistemic) do not seem reducible to evaluations of reliability is evidenced by the fact that we can intelligibly say (and observe others saying) things like 'he did the wrong thing (believed badly; is too irrational to be relied upon)', and so I'm not going to rely on him anymore, but I don't

to do with relationship-modification accounts in general. Wolf (2011, p. 336), for example, charges Scanlon's account of moral blame with being 'wimpy' or 'blame for wusses' because it fails to give sufficient attention to angry blame. If this is correct, then Smartt's real objection isn't that Boulton provides an unsatisfying account of *epistemic* blame, it's that relationship modification is not the right way to think about blame simpliciter.

¹⁰Boulton (2024, fn. 4) recognises this as a possibility in a footnote, but doesn't develop this thought.

blame him for it' (cf. Coates and Tognazzini, 2013, p. 9–10).¹¹ If Smartt is right that epistemic blame doesn't add anything to evaluations, then such locutions become rather puzzling, which is another observation that would need to be accounted for. Conversely, reporting that one blames another agent, having already identified them as irrational, seems to make one guilty of redundancy. Worse still, if – as I pointed out above – Smartt's point generalises to moral blame, then people who take themselves to be engaging in blame (understood as something distinct from a mere assessment of reliability) seem to be confused about what they are experiencing.

The potential redundancy in claims like 'He's irrational and I blame him for it' also suggests a way that Smartt may be hoisted by his own petard. Recall that the purpose of evaluation is to mark *different* kinds of sources according to *different* criteria. On Dogramaci's (2012) picture of evaluations, which Smartt endorses, different evaluations thus serve different functions. If we notice that there is a distinct kind of evaluation, then it is probably serving some kind of function that *isn't* being fulfilled by other kinds of evaluations. But this is precisely the kinds of grounds that should make us think that 'epistemic blameworthiness' is a distinct property which it is useful to keep track of!

Smartt's second objection, then, is unsuccessful against either Boulton's account or my own. However, his third argument is much more promising.

5. How many kinds of blame are there?

A final objection from Smartt is that allowing there to be a distinct, 'epistemic' type of blame opens the door to there being too many other kinds of blame. Here he charges Brown's approach with being too permissive: its structure allows for too many other kinds of blame that might stem from analogous belief–desire pairs, such as prudential blame (experienced when others violate prudential norms), or aesthetic blame in response to violations of aesthetic norms. Indeed, it can be hard to see why we should not allow all manner of different kinds of blame; if there can be 'office furniture norms', why not allow there to be 'office furniture blame' directed at office furniture 'norm violations'?

This objection is stronger than I think Smartt presents, because it could generalise to other accounts of epistemic blame. Boulton's account depends on there being a distinct kind of relationship between agents according to some normative ideal. It certainly seems that athletes on sports teams have a kind of relationship with one another as *teammates*, and they definitely blame one another for poor sporting performances. One can even blame one's rivals, if they break the rules, or even give up too easily. There could, in principle, be some kind of normatively guided aesthetic relationship between agents in the art world, for instance, comprising a set of intentions and expectations about how agents will behave with regard to some aesthetic standards. Perhaps some artists rely on one another to meet a certain set of aesthetic standards, and modification of this relationship on the basis of perceived faultiness in how one conducts themselves could be considered to be a form of blame? If Boulton takes epistemic blame to be concerned with a modification of trust that someone will live up to the standards of the epistemic relationship, why not allow agents to have a modification of trust that they will live up to other kinds of relationship standards too?

¹¹Boulton (2024) also argues there is an intuitive distinction between judging someone has fallen short of a standard, and being 'exercised or engaged' by this judgement. But rather than merely reporting different intuitions, I am trying to show Smartt fails by his own terms: his account actually leads to more unaccounted observations than the epistemic blame realist.

The agency-cultivation account is also not immune to this objection. I've previously argued that epistemic blame is justified by its ability to cultivate a distinctively valuable kind of epistemic agency. But we need not think that epistemic and moral agency are the only valuable varieties. It is valuable for people to be better athletes, artists, philosophers, chess players, dancers, politicians, etc. To be sure, one can get better at each of these by improving their epistemic agency, but improving epistemic agency alone is far from sufficient. Each of these fields has distinct standards, methods of improvement, cultures, roles, and internal set of goods, and agents can experience blame in any of them. Why not think there are also different types of blame present here then? If my account were also unable to answer this, Smartt's scepticism might succeed after all.¹²

If moral blame exclusively targets moral failings, it seems we need to either claim that blame in other domains is always unfitting (because not moral), that blame in other domains actually targets some hidden moral properties, or that there are distinctively non-moral kinds of blame. None of these options seem appealing. As we've seen, saying that blame is always unfitting outside of the moral domain would make our theory somewhat revisionist given our actual practices of blaming people who believe badly and warrant an explanation for why blame is only ever justified towards moral wrongs or bads. Saying that there are hidden moral reasons doesn't seem to quite fit the phenomenology of our blame, and some examples (e.g. the aesthetic domain) seem not particularly moral. But allowing for distinctive classes of blame, as the epistemic blame proponents would like, risks opening the door to all manner of types of blame.

I believe the third option has the most promise, because the unappealing consequence does not follow: we can allow for some non-moral types of blame without risking all manner of types to propagate. Here is how: there are moral norms, there are epistemic norms, but there is no distinct class of office furniture norms, or at least, not any that *matter*. 'Norms' in the sense we are concerned with are not merely potential kinds of reasons, or standards that one could theoretically point to relative to some domain. They are *authoritative* standards. They are the kind of thing that *we collectively recognise as important* and uphold as communities.

There is a tendency to sometimes think of epistemic norms as simply being what would characterise good reasoning and states ('don't hold inconsistent beliefs'), or to think of them in reference to some ideal theory, e.g. what a perfect Bayesian would believe. But this is not the relevant sense when it comes to thinking about blameworthiness and types of blame. Norms need to be the kind of thing that we normatively *expect* agents to follow, the failure of which licenses negative responses given the circumstances that the agent was in. An analogy will help. Suppose that consequentialism is, in fact, the correct normative theory. This alone does not entail that agents who fail to maximise good consequences are thereby blameworthy. Consequentialists regularly made a distinction between what is wrong and what is blameworthy as a defence against charges of demandingness (e.g. Morris, 2017). Returning to the epistemic domain, we need to take care to not move from noticing that a certain belief or inference violates some epistemic standard without an excuse to inferring that the agent is thereby an appropriate target of blame. Their violation needs to be the kind of thing that we expect people to live up to, and which reflects *badly* on them as an agent, not merely reflects on them as a less than perfect being.

¹²The strength of this objection may face mixed intuitions. It is worth noting that some philosophers *do* think that there are such things as aesthetic responsibility and blame (Nelkin, 2020; Kubala, [forthcoming](#); Wolf, 2016), and that we do have prudential duties towards ourselves. Additionally, the blame we experience in domains that are not clearly moral or epistemic needs to be accounted for; sceptics cannot say that the realists have failed to do a sufficient job if they cannot account for this either.

Returning to Smartt's objection: it is not enough to merely observe that there could be some other set of non-moral and non-epistemic standards, or that some people think that there are some sets of standards which people can fail to comply with. One needs to show that these are the kinds of things that we expect people to recognise and treat as authoritative and think reflect badly on them independent of any moral or epistemic considerations that warrant experiencing the kinds of reactions we associate with blame.

One might ask why moral and epistemic norms have this authority, but others do not. It is worth noting that this question is essentially asking about the foundations of normativity, and it seems an unfair standard to insist that epistemic blame realism should be considered tenable only if realists have first solved *that* issue. Admitting epistemic norms have a *bona fide* authority (Grimm, 2009) does not compel us to admit that other kinds of merely possible norms thereby are actually authoritative too, and thus nor are we compelled to say that there are other kinds of blame.

Appealing to distinct types of norms doesn't particularly help Brown. Given she justifies her account with reference to what blame is (a belief-desire pair), the role that epistemic blame plays in some epistemological debates, and instances of people blaming others for believing badly, there is nothing here to rule out people having other kinds of belief-desire pairs and blaming others for other kinds of violations. Nor does it seem to help Boulton: there is no inherent feature of moral/epistemic relationships and moral/epistemic norms as he understands them to limit other kinds of relationships allowing for other kinds of blame. But this answer *does* dovetail particularly well with the agency-cultivation account in Piovarchy (2021). Though I argued there that agency cultivation is the goal of blame, this is not justified for the agent's sake, but for the sake of the community. In later work, I've been more explicit about this (Piovarchy, 2023; *ms*), and clarified that I now think it's more accurate to think of blame's function as upholding valuable norms. This is a minor extension: while I originally argued that we blame people to cultivate their agency, the kind of agency we want is one that is responsive to authoritative moral and epistemic considerations; that is, one which complies with moral and epistemic *norms*.

This story is compatible with what is called social epistemic instrumentalism, which has been developed by others in a number of places (Dyke, 2021; Fleisher, 2025; Hannon & Woodard, *forthcoming*; cf. Chrisman, 2022; Boulton, *forthcoming*). The basic idea is that moral and epistemic norms are particularly important for the functioning of communities, given how reliant we are on each other for our well-being. We need people to treat us well, and to be reliable sources of information, and agreed-upon standards help maintain the systems that provide these benefits. When we blame someone, we are saying that they failed to comply with the norm, and that they need to improve their behaviour. We are signalling that the norm matters, and communicating that we will not tolerate continued violations, thereby providing an incentive for them to comply, and helping them learn what the contours of the norm are. Once they have become competent with the norm and internalised it, they are then able to help maintain valuable forms of both moral and epistemic co-operation (cf. Henderson & Graham, 2019). They not only treat others well and avoid engaging in destructive behaviour, they also act as a reliable source and conduit of information, whether that be on matters they have first-hand experience with or regarding information they have heard from others.

Note this explanation makes sense of our ambivalence or disagreement about aesthetic and athletic blame: our uncertainty stems from whether we take the norms in those fields to be authoritative. Sports teams need to uphold certain standards in order to function well. If athletes and fans think that sporting excellence is a genuinely valuable activity worth cultivating among professional athletes, it makes sense that the former would blame poor performers who could reasonably be expected to do better. In contrast, while many artists care about aesthetic standards and expect others to uphold

them, philistines like myself can find it difficult to see what the fuss is about and accordingly won't take such norms to be authoritative.¹³

This does not provide a full account of when norms are valuable or authoritative, or under what conditions other sets of norms that we find alien could come to have genuine importance. But epistemic blame proponents do not need a comprehensive account of the foundations of normativity in order for epistemic blame realism to be viable, any more than moral blame theorists need to first demonstrate that moral realism is true. It is enough to have made clear what the possibility of there being types of blame will depend upon, and provided a plausible answer at that. Blame functions to uphold norms, moral and epistemic norms are genuinely authoritative, so there are distinctly moral and epistemic forms of blame.¹⁴ Other merely potential types of norms don't serve any especially valuable function in communities, nor are they recognised to, so it makes sense that there won't be corresponding types of blame. The connection between norms and blame, however, only fits well with the agency-cultivation account, which also happens to be the only account of epistemic blame capable of avoiding each of Smartt's objections.¹⁵ The epistemic blame realist, then, lives to epistemically blame another day.

6. Conclusion

At this point, I think it would help to take stock of the dialectic between realists and sceptics, to see why certain moves may be unlikely to sway the other, which the arguments from Matheson and Milam and Smartt help illuminate.

We observe that people blame others for their epistemic failings and want to theorise about these interactions accurately. For any given instance in which one group of people – the epistemic blame realists – claim there is a distinct kind of blame that targets distinctly epistemic failings, their opponents – epistemic blame sceptics – can claim that either the blame seems intuitively fitting, but is moral in nature, or else is unfitting. For any epistemic failing that sceptics claim realists ought to blame (or ought to believe are blameworthy), and yet don't, the realist can claim that though blame would be fitting, this happens to be a case in which we don't care.¹⁶ Appeals to intuitions about such cases

¹³This is not to say that the experience is wholly alien; Susan Wolf once remarked in a lecture that Star Wars fans certainly blame George Lucas for his aesthetic choices in the prequels.

¹⁴This provides a reply to Matheson and Milam's earlier arguments that blaming others for non-moral norm violations doesn't provide good consequences, and that we never have epistemic obligations. Though they don't explain how they understand 'obligation', what matters is that social epistemic instrumentalism provides an account of the authority of epistemic norms and the appropriateness of blaming culpable violations.

¹⁵Smartt also makes an objection that Boul's account does not have the resources to distinguish between intuitively blameworthy and excused agents, since excused agents are also unreliable and thus seem to impair the epistemic relationship with others when they believe badly. I have omitted engagement with this objection for a few reasons. First, my own account, being explicitly capacitarian, can easily distinguish between excused and blameworthy agents—agents are excused when they lacked the capacity to believe the right thing for the right reasons. Conveniently, capacities come in degrees, so this also explains how blameworthiness can also be mitigated. Second, Boul (2024) has replied to this point elsewhere, and getting into the weeds of his adoption of Scanlon's account of meaning and significance would take us too far afield. I believe Boul's reply is unsuccessful. For a thorough examination of why, see Piovarchy (forthcoming-b).

¹⁶This dialectic helps show why another objection from Smartt—that Brown's theory can't quite account for the force of blame—may also not convince everyone. His worry is that if the desire that the agent believed better must have a certain level of strength (otherwise the blame will have no force), but we simply do not experience desires with this strength, which he supports with some examples that will not sway realists. When Smartt claims that he doesn't think his desire that Caligula lived a better life amounts to blame, the realist will reply that, so long as we are clear that it is a desire that Caligula had lived a *moral* life, because of

risk failing to advance the argument between sceptics and realists in either direction. Appeals to one party having the burden-of-proof also don't seem especially obvious given our tendency to blame one another for non-moral norm violations without invoking moral considerations, and given most people are bad at generating philosophical arguments on the foundations of normativity.

Smartt tries to bolster the sceptical argument and break the stalemate by focusing on purported differences between moral and epistemic blame. He argues that moral blame is much stronger, made in response to victims being harmed, and typically signals concern for said victims. The realist can account for these differences by saying that we care less about norm violations that do not feature victims, and that blame need not necessarily signal concern for victims. The realist can further weaken Smartt's approach by pointing out they are not making any special pleading, as many moral norm violations also do not feature victims and are accompanied by weaker blame, which also does not signal any concern for victims.

While these objections to epistemic blame in general don't resolve the debate, objections to particular accounts of epistemic blame seem more successful. Brown's desire that others not have believed badly has a somewhat tenuous connection to what we think of as blame, and Boulton's account seems to lack the resources to distinguish between intuitively blameworthy and blameless agents (cf. Piovarchy, [forthcoming-b](#)). Both Boulton and Brown are at risk of allowing too many other types of blame, provided we can find a relevant belief-desire pair or some kind of normative ideal capable of grounding a relationship. However, the account originally put forward in Piovarchy (2021) – notably an account that explicitly provides an explanation of what epistemic blame is, when epistemic blame is warranted, and what justifies our practices of blaming – avoids all of these objections.¹⁷ While new objections may arise in the future, for now it seems that proponents are entitled to continue blaming others for their purely epistemic faults. At least, they have good evidence for believing this is the case and won't be blameworthy if they turn out to be mistaken.

our concern for *morality*, and our negative reactions express this, then this seems like a very plausible characterisation of blame. When Smartt claims that he does not blame Lewis for holding inconsistent beliefs, the realist will claim that this is simply a way of saying that we do not blame every single epistemic norm violation we are aware of. Smartt finally charges Brown with a kind of inconsistency, as epistemic blame is supposed to be a systematic feature of our practices, but we rarely engage in epistemic blame. But it isn't clear that epistemic blame is rare. Though *purely* epistemic blame may be rare, realists accept that many instances of blame is both moral and epistemic in nature (e.g. Brown, 2020, p. 15), given we often have moral reasons to comply with epistemic norms. Smartt seems to be thinking that epistemic blame is relatively rare because a high proportion of epistemic norm violations do not generate epistemic blame, but this is not the same as showing that epistemic blame is uncommon in the sense of not being encountered often, or not being a regular part of our practices.

¹⁷An anonymous reviewer wonders how agents could ever be epistemically blameworthy given the following: suppose an agent seems to have the capacity to engage in some doxastic performance. They instead exhibit a subpar doxastic performance. Something must have prevented them from engaging in the superior performance. But unless they are responsible for that 'something', it seems like in fact they lacked the capacity to engage in the better doxastic performance. A full reply relies on understanding capacitarian and reasons-responsive approaches to blameworthiness, which are quite well-established in the moral responsibility literature. The basic answer is that it matters greatly what we 'hold fixed' when looking at relevant counterfactuals or nearby possible worlds, and agents can retain the (responsibility-relevant) capacity to do/believe the right thing even if in some worlds they fail. For some influential treatments of these issues, see Fischer and Ravizza (1998), Brink and Nelkin (2013), Jaster (2020). For an application of this approach to culpable ignorance, which will generalize to cases of epistemically bad believing, see Piovarchy (2022).

Author Biography. Adam Piovarchy is a Research Fellow at The University of Notre Dame, Australia. His work focuses on moral responsibility, blame, and hypocrisy.

References

- Ballarini C.** (2022). 'Epistemic Blame and the New Evil Demon Problem.' *Philosophical Studies* **179**(8), 2475–505.
- Boult C.** (2021a). 'Standing to Epistemically Blame.' *Synthese* **199**(3–4), 11355–75.
- Boult C.** (2021b). 'There is a distinctively epistemic kind of blame.' *Philosophy and Phenomenological Research* **103**(3), 518–34.
- Boult C.** (2023). 'The Significance of Epistemic Blame.' *Erkenntnis*.
- Boult C.** (2024). 'Epistemic blame as relationship modification: reply to Smartt.' *Philosophical Studies* **181**(2), 387–96.
- Boult C.** (forthcoming). 'The Relational Foundations of Epistemic Normativity.' *Philosophical Issues*.
- Brink D. O., & Nelkin D. K.** (2013). 'Fairness and the architecture of responsibility.' In *Oxford Studies in Agency and Responsibility* (Vol. 1), pp. 284–313. Oxford University Press.
- Brown J.** (2017). 'Blame and Wrongdoing.' *Episteme* **14**(3), 275–96.
- Brown J.** (2020). 'Epistemically Blameworthy Belief.' *Philosophical Studies* **177**(12), 3595–614.
- Chrisman M.** (2022). *Belief, Agency, and Knowledge: Essays on Epistemic Normativity*. Oxford University Press.
- Coates D. J. & Tognazzini N. A.** (2013). 'The Contours of Blame.' In D. Justin Coates & Neal A. Tognazzini (eds), *Blame: Its Nature and Norms*. Oxford University Press. pp. 3–26.
- Dogramaci, S.** (2012). 'Reverse Engineering Epistemic Evaluations.' *Philosophy and Phenomenological Research* **84**(3), 513–30.
- Dyke M.M.** (2021). 'Could our Epistemic Reasons be Collective Practical Reasons?' *Noûs* **55**(4), 842–62.
- Fischer J. M., & Ravizza M.** (1998). *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge University Press.
- Fleisher W.** (2025). 'Epistemic Practices: A Unified Account of Epistemic and Zetetic Normativity.' *Noûs* **59**(1), 289–314.
- Grimm S. R.** (2009). 'Epistemic Normativity.' In A. Haddock, A. Millar, and D. Pritchard (eds), *Epistemic Value*, pp. 243–64. Oxford: Oxford University Press.
- Hannon, M. & Woodard, E.**, (forthcoming). 'The Construction of Epistemic Normativity.' *Philosophical Issues*.
- Henderson D. & Graham P.** (2019). 'Epistemic Norms as Social Norms.' In M. Fricker, N. J. L. Pedersen, D. Henderson and P. J. Graham (eds), *Routledge Handbook of Social Epistemology*, pp. 425–36. Routledge.
- Kauppinen A.** (2018). 'Epistemic Norms and Epistemic Accountability.' *Philosophers' Imprint* **18**, 1–16.
- Kubala R.** (forthcoming). 'Aesthetic Blame.' *Journal of the American Philosophical Association*.
- Matheson B. & Milam P.** (2022). 'The Case Against Non-Moral Blame. In M.C. Timmons (ed), *Oxford Studies in Normative Ethics, Volume 11*, pp. 199–222.
- Morris R.** (2017). 'Praise, Blame, and Demandingness.' *Philosophical Studies* **174**(7), 1857–69.
- Nelkin D. K.** (2020). 'Equal Opportunity: A Unifying Framework for Moral, Aesthetic, and Epistemic Responsibility.' *Proceedings of the Aristotelian Society* **120**(2), 203–35.
- Piovarchy A.** (ms). Signalling, Sanctioning and Sensitising: How to Uphold Norms with Blame.
- Piovarchy A.** (2021). 'What do We Want from a Theory of Epistemic Blame?' *Australasian Journal of Philosophy* **99**(4), 791–805.
- Piovarchy A.** (2022). Situationism, Capacities and Culpability. *Philosophical Studies*. **179**(6), 1997–2027.
- Piovarchy A.** (2023). 'Vice epistemology, norm-maintenance and epistemic evasiveness.' *Synthese* **201**(105), 1–20.
- Piovarchy A.** (forthcoming-a). 'Epistemic hypocrisy and standing to blame.' *Erkenntnis* 1–21.
- Piovarchy A.** (forthcoming-b). 'Epistemic Blame Isn't Relationship Modification.' *The Philosophical Quarterly*.
- Rawls J.** (1955). 'Two Concepts of Rules.' *The Philosophical Review* **64**(1), 3–32.
- Rettler L.** (2018). 'In Defense of Doxastic Blame.' *Synthese* **195**(5), 2205–26.
- Scanlon T.** (2008). *Moral Dimensions: Permissibility, Meaning, Blame*. Cambridge, Massachusetts: Belknap Press of Harvard University Press.
- Schmidt S.** (2021). 'Epistemic Blame and the Normativity of Evidence.' *Erkenntnis* 1–24.

- Smartt T.** (2024). 'Scepticism about Epistemic Blame.' *Philosophical Studies*.
- Vargas M. R.** (2013). *Building Better Beings: A Theory of Moral Responsibility*. New York: Oxford University Press.
- Williamson T.** (forthcoming). 'Justifications, Excuses, and Sceptical Scenarios.' In F. Dorsch and J. Dutant (eds), *The New Evil Demon*. Oxford: Oxford University Press.
- Wolf S.** (2016). 'Aesthetic Responsibility.' *The Amherst Lecture in Philosophy* **11**, 1–25.