

# The Epistemic Norm of Blame

D. Justin Coates<sup>1</sup>

Accepted: 20 September 2015 / Published online: 30 September 2015  
© Springer Science+Business Media Dordrecht 2015

**Abstract** In this paper I argue that it is inappropriate for us to blame others if it is not reasonable for us to believe that they are morally responsible for their actions. The argument for this claim relies on two controversial claims: first, that assertion is governed by the epistemic norm of reasonable belief, and second, that the epistemic norm of implicatures is relevantly similar to the norm of assertion. I defend these claims, and I conclude by briefly suggesting how this putative norm of blame can serve as the basis for general norms of interpersonal generosity.

**Keywords** Blame · Epistemic norms · Assertion

## 1 Introduction

It is widely agreed that it is inappropriate to hold an agent responsible for an action if she is not morally responsible for performing that action. After all, if she is not morally responsible for that action, then she cannot be *blameworthy*—i.e., *deserving* of blame. And so, it cannot be appropriate for us to hold her responsible for that action. But of course, even if an agent genuinely deserves to be blamed for an action, it doesn't thereby follow that it is appropriate for us to hold her responsible for that action. For example, as Gary Watson (1987), G. A. Cohen (2006), Angela Smith (2007), and others have pointed out, even in cases in which wrongdoers deserve blame, *we* might lack the moral standing to blame. Or alternatively, if another's moral transgression is too minor, or if the wrongdoing is simply none of our business, or if we are complicit in the wrongdoing, or if it would be hypocritical for us to blame, then quite plausibly, it is not all things considered appropriate for us to blame.<sup>1</sup>

---

<sup>1</sup>See also, T. M. Scanlon (2008), R. Jay Wallace (2010), Patrick Todd (2012), and D. Justin Coates and Neal Tognazzini (2012, 2013) for more on the conditions under which it is all things considered appropriate to blame.

✉ D. Justin Coates  
djcoates@uh.edu

<sup>1</sup> Department of Philosophy, University of Houston, Houston, TX, USA

A full “ethics of blame” will take all of these considerations (and more) into account, telling us when it is, and when it is not, appropriate to blame.<sup>2</sup> In this paper, then, I will take one small step in articulating such an ethics. In particular, I’ll consider one of the norms that govern blame, one that has heretofore been unremarked upon (though much has been made of its status as a norm of other social practices)—the epistemic norm of blame.

The main contention of this paper is that the activity of blame is answerable to a distinctively epistemic norm, and that this fact has significant practical implications for the practices of blaming and holding others morally responsible.<sup>3</sup> The norm in question is ‘epistemic’ in the sense that an agent will satisfy it only if she meets some (to be described) epistemic standard. Of course, this means that the norm might ultimately be connected to a more general epistemic norm of practical reasoning or action, but seeing how such a norm applies to our blaming practices is important, since blame (and its expressions) can be coercive and can cause harm and suffering. More specifically, then, I want to argue that the following norm governs our practice of blaming:

#### Epistemic Norm of Blame (ENB)

It is inappropriate (absent special justification) for *A* to blame *B* for *x*-ing if it is not reasonable for *A* to believe that *B* is morally responsible for *x*-ing.<sup>4</sup>

As stated ENB is a condition on justified instances of blame such that, even if *B* is in fact morally responsible for *x*-ing, it may still be inappropriate for *A* to hold *B* responsible. In particular, this will be true in cases in which *A* has failed to acquire evidence of the sort that renders her beliefs about *B*’s blameworthiness epistemically reasonable or justified, or in cases in which *A* has formed her belief concerning *B*’s status as morally responsible from an unreliable or improperly functioning belief-forming mechanism.<sup>5</sup>

Taken at face value, the idea that blame is governed by ENB might seem incredibly far-fetched on the one hand, or almost trivial on the other. It might seem incredible, after all, because if an agent is in fact morally responsible for an action, why should the facts about blamers’ epistemic status matter? Indeed, if it were true that ENB governs our blaming practices, it would seemingly offer wrongdoers a legitimate basis upon which to skirt

<sup>2</sup> By discussing an “ethics of blame,” I do not mean to suggest that all of these considerations are unique to the norms of blaming. For example, it might be that just as our lack of a standing relationship precludes me from being warranted in blaming you (at least in some cases), so too that lack of a standing relationship would preclude me from being in asking a serious favor of you as well. Thus, a full ethics of blame takes into account all of the norms that might govern blame, even if in many cases, such norms are not internal to the activity of blame itself.

<sup>3</sup> This claim will come as no surprise to those who think that action itself is governed by an epistemic norm, since although we sometimes blame others in virtue of our attitude (s) towards them, expressions of blame are typically actions (e.g., rolling one’s eyes, telling another off, aggressively honking the car’s horn, etc.). The classic statement of this view is due to John Hawthorne and Jason Stanley (2008). As you’ll no doubt notice, however, I accept a more permissive epistemic norm than the one defended by Hawthorne and Stanley.

<sup>4</sup> In saying that *A* blame is inappropriate if it is not reasonable for her to believe that *B* is morally responsible, I mean to pick out a notion of “reasonability” that is closely allied with epistemic justification. So “not reasonable” could be filled out as “unreasonable,” “unjustified,” or “unwarranted.” Following Jennifer Lackey (2007), I prefer framing this norm in terms of what it would be *reasonable* for *A* to believe, but I have no principled reason for thinking that ENB couldn’t be framed in terms of other epistemic goods.

<sup>5</sup> I want to make clear here that ENB is meant to be compatible with multiple theories of which properties a belief must instantiate in order to be epistemically reasonable. I have noted three candidates here (evidence, reliability, and proper functioning), but there are no doubt others. For the purposes of this paper, I will be agnostic about the precise conditions under which it is or is not reasonable to believe that an agent is morally responsible for her action in the hopes of achieving an ecumenical consensus concerning the fact that the activity of blame is governed by an epistemic norm. Accordingly, if you’re an internalist, then fill in the content of “reasonable for *A* to believe that ...” in your preferred way, and if you’re an externalist, fill it in in your preferred way.

responsibility. However, “look, you don’t have any good evidence that I’m responsible…” is hardly a compelling defense against blame in most contexts. Unlike “I didn’t do it,” “what I did wasn’t wrong,” or even “that’s rich coming from *you*,” which are defenses that, if true, clearly seem to undermine the legitimacy of blame, appealing to a blamer’s putative violation of ENB is different: quite simply, it falls remarkably flat. After all, in those cases in which you are guilty, you are aware of this, and so pointing to blamers’ violations of ENB would appear only to be a facile attempt to escape the blame that you rightly deserve.<sup>6</sup> ENB, then, seems somewhat suspect. Maybe a related epistemic norm protects procedural rights in legal contexts, where so much is at stake, but in interpersonal cases of blame and informal sanctioning, ENB appears to be overly restrictive at best, the last refuge of scoundrels at worst.

On the other hand, however, ENB can seem to be almost a truism or platitude. *Of course*, you might think, ENB (or some ENB-like principle) governs our blaming practices. It’s just not fair to go around willy-nilly blaming, criticizing, and condemning others with no reason to think they’re deserving of such treatment, even though such blame, criticism, and condemnation will, as a matter of fact, sometimes (frequently, even) hit upon targets deserving of these forms of treatment. But let’s suppose this is correct and (briefly, for the sake of argument) simply grant the truth of ENB. It doesn’t thereby follow that we have an explanation of *why* ENB is true, since the (putative) obviousness of the truth of ENB won’t itself entail an explanation as to why ENB is true. Murdering for convenience, for example, is also obviously wrong, even though explaining why this is the case is notoriously difficult.<sup>7</sup> Consequently, even granting that ENB has the status of a commonplace or platitude, we still need a better understanding of precisely *why* it is that we’re obligated to blame only if we meet some epistemic standard. In other words, even if its status as a genuine norm could be taken for granted in discussions of the ethics of blame, there’s still more to say in defense of ENB. Fortunately, I think we can muster a compelling defense of ENB, which explains its normative basis and thereby allays any suspicions one might have about ENB, while also supplementing our understanding of the norm and of its practical implications.

## 2 The Argument from Implicature

On behalf of ENB, I offer the Argument from Implicature (AFI). AFI goes like this:

1. It is inappropriate (absent special justification) to assert that  $p$  if it is not reasonable for you to believe that  $p$ .
2. The epistemic norm (s) of assertion are relevantly similar to the epistemic norm (s) of implicature.
3. Thus, it is inappropriate (absent special justification) to implicate that  $p$  if it is not reasonable for you to believe that  $p$ .

<sup>6</sup> These considerations do not seem to arise in quite the same way in the case of a third party coming to someone’s defense. Indeed, it often helps us calm down when others remind us that we don’t know all the facts and that we shouldn’t jump to conclusions. Of course, if the third party knows that the agent in question is guilty, it would be disingenuous for her to suggest otherwise or appeal to ENB as a reason to refrain from blaming. But this still leaves open that some third-party appeals to ENB might be normatively significant in a way that first-person appeals typically are not.

<sup>7</sup> For a classic discussion of this difficulty, see Barbara Herman (1989).

4. But when *A* blames *B* for *a*-ing, *A* implicates, among other things, that *B* is responsible for *a*-ing.
5. Therefore, ENB: it is inappropriate (absent special justification) for *A* to blame *B* for *a*-ing if it is not reasonable for *A* to believe that *B* is responsible for *a*-ing.

As it stands, AFI is simple and straightforward, and I think it provides us with a powerful explanation of ENB. Of course, as is the case with almost any interesting argument, AFI relies on premises that are controversial. For example, it's not altogether clear in what sense it might be inappropriate for someone to assert that *p* without being reasonable in his or her belief that *p*. Is this merely, for example, a conversational norm? And if so, why should we care about violations of ENB? So too, why should we think that assertion and implicatures really are governed by the same set (or relevantly similar sets) of epistemic norms? Indeed, you might even worry about the claim that assertion and implicature are governed by an epistemic norm at all.

These are important questions that the defender of AFI needs to deal with. So in what follows, I will provide a more systematic defense of AFI's premises and its conclusion.

### 3 Assertion and Reasonable Belief

First, AFI relies on a claim about the epistemic norm of assertion, and this claim is sure to be controversial. But although these norms have been widely discussed as *constitutive* norms of assertion—that is, norms that individuate assertion from other speech acts—and this is a matter of some controversy, as I've stated it, (1) need not (but could) be a constitutive norm of assertions *qua* speech acts.<sup>8</sup> That is, even if assertions are not identified by a constitutive epistemic norm, (1) might still be true, since though failing to conform to (1) would not be a failure *qua* assertion, independently, it seems to be a *moral* failing. After all, if I assert to you that *p*, then *ceteris paribus*, you can now be rational in acting on *p* (or more minimally, you can now be rational in using that *p* as a premise in your practical reasoning). And this can, in many cases, lead to very bad consequences for you. It is thus morally problematic for me to put you in such a situation without sufficient justification (which I obviously lack when I assert that *p* without being reasonable in my belief that *p*). In other words, the rationality of your acting on *p*, where *p* is something you come to believe as a result of my assertion, depends (in large part) on the reasons I have for believing that *p*. And if you act on *p* even though I'm not reasonable in my belief that *p*, then you could act irrationally as a result. And by acting irrationally, you might very easily find yourself in an injurious or otherwise bad situation. Thus, I shouldn't purport to give you a reason to act on *p* if I am not reasonable in my belief that there is in fact a reason to act on *p*, since doing so can range from reckless to mean-spirited to (in extreme cases) evil.

The above argument might be somewhat opaque, so to see an illustration of precisely why it is that we shouldn't (absent special justification) assert that *p* without being reasonable in our belief that *p*, consider the following case.

<sup>8</sup> For more on the claim that something like (1) is a constitutive norm of assertions *qua* speech acts, see Timothy Williamson (1996) and Lackey (2007).

### 3.1 Wrigley Field

Steven is visiting Chicago for the first time, and he wants to see a game at Wrigley Field. While walking in the Loop, he asks Sylvia, who looks as though she might be a local, for directions to Wrigley. Although Sylvia remembers that the field is on Addison, she can't remember whether Steven should take the Blue Line to Addison or the Red Line to Addison. So not wanting to appear ignorant, but without really having any reason to think that it is the Red Line rather than the Blue Line, Sylvia simply asserts, "You can get there by taking the Blue Line to Addison. Once you get off, the field will be right in front of you."

Unfortunately for Steven, Sylvia gets it wrong—to get to Wrigley Field from the Loop, you take the Red Line to Addison (or some combination of the Brown and Red Lines). So by taking her assertion that he could get to Wrigley Field by taking the Blue Line to Addison to heart, Steven ends up somewhere to which he had no reason to go.

If you're like me, you're probably thinking that Sylvia's not just violating conversational norms in *Wrigley Field*, but that she's also done something that is morally objectionable. The thought here is that the problem with Sylvia's assertion isn't simply one that you would have from the point of view of a conversational interlocutor—no different, for example, than if Sylvia had failed to satisfy the Gricean maxim of relation (i.e., the maxim: *be relevant*).<sup>9</sup> There is also (or perhaps, instead) a moral failure on Sylvia's part. After all, rather than simply saying "I'm not sure" or presenting Steven with her evidence and letting him decide, she presented unknown information as information for which she had a reasonable basis. Consequently, her testimony gives Steven *prima facie* license to buy a ticket and to wait for a Blue Line train (instead of a Red Line train to the ballpark). Indeed, her assertion seems to give Steven reasons to do these and all sorts of other things that are in fact counter to his intentions and goals, since the Blue Line intersects Addison three miles to the west of where the Red Line intersects Addison. And if he acts on these putative reasons, then it seems that, without any justification, Sylvia has, at the very least, inconvenienced Steven a great deal.

And look, in general, it's morally *wrong* to mislead others in this way, to impede their intentions or goals, or to otherwise significantly inconvenience them without sufficient justification, even if you do so unintentionally. But the real issue isn't just that Sylvia misleads Steven in this case. Rather, it is the lack of regard she has for him—a lack of regard that manifests itself in her willingness to harm him by being misleading in this case. However, even if Steven is not ultimately misled—if, for example, he asks someone else, and they give him the right instructions—Sylvia has still wronged Steven by engaging with him in a manner that shows undue concern for him, for his interests, and for the truth.<sup>10</sup>

In general, this means that we shouldn't assert what we aren't reasonable in believing because when we do so, (i) we unduly and negligently risk significant harm to others, and (ii) we show insufficient concern for others and their interests. Moreover, in many cases, the potential harm is much greater than the relative minor inconveniences that Steven will experience when he gets of the Blue Line at Addison and discovers that he's wasted a couple

<sup>9</sup> I suspect that conversational norms are broader than those just the set of norms considered by Paul Grice, but his conversational maxims are a very good attempt at circumscribing these norms. See Grice (1975).

<sup>10</sup> It's worth pointing out the criticism of Sylvia here is closely connected to Harry Frankfurt's criticism of bullshitters, since the bullshitter, unlike the liar, is wholly unconcerned with the truth. Similarly, Sylvia seems more concerned with presenting herself as knowledgeable rather than getting it right for Steven. For more, see Frankfurt (1986).

of hours and maybe even missed the game.<sup>11</sup> Accordingly, it seems that we have a weighty reason to refrain from asserting propositions with respect to which we are not reasonable in believing. Hence, Sylvia's wrongdoing in *Wrigley Field* helps us see that, quite plausibly, (1) is true.

### 3.2 ENB and Assertoric Blame

Before going on to defend the other premises of AFI, I think it's worth pointing out here that if (1) is true then, it will directly follow that many instances of blame will be governed by ENB. For in many (but certainly not all) cases of *expressed* blame, we are able to blame another for their action (s) simply by asserting something about the agent in question: "You're a *jerk*," or "that was a shoddy thing to do," for example. Of course, not all assertions of the type "you're a jerk!" are expressions of blame, as is the case when one friend jokingly tells the other that she was a jerk, without thereby meaning to blame or condemn her friend for that. (Indeed, it is sometimes possible to express good will and intimacy with the very assertions that, in other contexts, express blame and condemnation.) Similarly, not all utterances of "you're a jerk" count as assertions—some such utterances are purely expressive of the utterer's attitudes, and in no way mean to represent the world.<sup>12</sup> But in cases in which a speaker asserts that someone is a jerk, or that she is morally responsible for some bad outcome, or that she has acted wrongly, or that he is a bad man, etc., with the aim of condemnation, it seems possible that in such a context, the speaker has thereby *blamed* the jerk, the responsible agent, the wrongful actor, the bad person.

If we blame another agent by asserting something about that agent and we also lack good reasons for believing him or her to be morally responsible, then we risk misleading others about this agent. That's bad enough, of course, but it's not what's really awful about blame assertions. After all, it's not always the case that we risk misleading people—perhaps no one hears us, or no one who does hear us will be inclined to believe us. It still seems inappropriate to blame (absent good reasons for believing the would-be blamee to be morally responsible). This is because, I think, most of us attach some significance to be value of innocence and also to the value of *just deserts*. What I mean by this is that it is very plausible that we shouldn't blame people if they do not deserve to be blamed. (Indeed, by my lights, this principle is at the heart of interpersonal moral engagement.) Thus, when we blame someone without satisfying ENB, we reveal ourselves to be insufficiently concerned with the value of innocence or with making sure that we blame only if the agent in question deserves to be blamed. And it's *this* moral failing, and not merely the willingness to mislead, which is the real problem with blame assertions that fail to satisfy ENB.<sup>13</sup>

<sup>11</sup> You see this point being emphasized as the basis for due process in the criminal law, since there is significant harm done to those who are wrongly accused of crimes or who are incorrectly found to be guilty of crimes that they did not commit. Justice Brennan (1970), in his majority opinion in the decision of *In re Winship* emphasizes this: "The reasonable doubt standard plays a vital role in the American scheme of criminal procedure. It is a prime instrument for reducing the risk of convictions resting on factual error ... The accused, during a criminal prosecution, has at stake interests of immense importance, both because of the possibility that he may lose his liberty upon conviction and because of the certainty that he would be stigmatized by the conviction. Accordingly, a society that values the good name and freedom of every individual should not condemn a man for commission of a crime when there is reasonable doubt about his guilt."

<sup>12</sup> This seems less true in the case of utterances of the type "that was a shoddy thing to do." But even if blame itself is never assertoric, it seems plausible to think that in blaming you I have expressed my blame, I am implicating the truth of your being blameworthy. I'll return to cases in which I merely implicate your blameworthiness in II.3.

<sup>13</sup> As an anonymous referee has rightly pointed out to me, framing things in this way very much pits the value of innocence against the value of blame. For more on how we should understand this latter value, see Franklin (2013).

Thus in these sorts of cases, which do not seem to be altogether rare, an assertion can function to blame. And if we think that it's important to refrain from blaming people that are innocent (or more generally, that do not deserve to be blamed), then we should not be willing to do so in circumstances in which there are not good grounds for believing the actor in question to be deserving of blame. Consequently, if (1) is true, then these instances of blame are governed by ENB *even if* upon reflection, we determine that (2) is false. However, it seems to me that if (2) is true, so I turn to a defense of this premise of AFI now.

#### 4 Assertions and Implicatures

If, as I have argued above, AFI's (1) is truly a normative constraint on our assertions, the soundness of the argument becomes much easier to defend. For plausibly, if it is inappropriate (absent special justification) for me to assert that  $p$  if I do not have good reason to believe that  $p$ , then how could it be appropriate for me to implicate (either through conversational or conventional forms of implicature) that  $p$  if I do not have good reason to believe that  $p$ ? This would be surprising given that assertions and implications often have the same conversational functions. Putting this point a bit more concretely: if I shouldn't assert that Smith is a bad student because I don't have any reason to believe this, then why would it be okay, given the exact same epistemic state with respect to Smith's qualities as a student, for me to implicate that she is a bad student by going on and on about her punctuality and excellent penmanship while ignoring completely the quality of her scholarship?

Furthermore, let us suppose that rather than asserting that Wrigley Field was off of the Blue Line, Sylvia had said, "it's definitely not the Red Line" after having explicitly been confronted with the options. Or imagine that she said, "You'll pass Wrigley Field when you take the CTA to O'Hare," after having already discussed with Steven that you take the Blue Line to O'Hare Airport. In either case, Sylvia clearly implies that Steven can get to Wrigley Field by taking the Blue Line to Addison. Of course, like her assertion that he could get to Wrigley Field by taking the Blue Line, Sylvia's implicature also leads Steven to get on a Blue Line train, ride north for at least 25 min, etc. Consequently, it seems that the assertion "it's on the Blue Line" and the implication that Wrigley is on the Blue Line (which was performed by the utterance "it's not on the Red Line" in a conversational circumstance where those were the only two options or by the utterance "you'll pass it on the way to the airport" in conversation that presupposes that one only gets to the airport by taking the Blue Line) play the exact same conversational role. And given this, we should expect assertion and implicatures to be governed by relevantly similar norms.

Moreover, if *asserting* that Wrigley Field is on the Blue Line is wrong (in part) *because* it risks undue harm for Steven, it seems that *implicating* that Wrigley is located at Addison on the Blue Line would be similarly wrong because such implicature risks harming Steven in *exactly the same way*. And if this is correct, then just as (2) tells us, we should conclude that the normative standards on assertion are relevantly similar to the normative standards on implicature. From (1) and (2), then, it straightforwardly follows that (3).

#### 5 Implicatures and Blame

Having shown (3) to be plausible, I think it's easy to make the case that (5)—i.e., ENB—is also plausible. For (4), which links (3) and (5) seems undeniable. To see this, consider the

widely accepted thesis concerning the relationship between appropriate blame and moral responsibility that I began with: the *Desert Thesis*. According to this *Desert Thesis*, *A* is permitted to blame *B* for *a*-ing only if *B* is morally responsible for *a*-ing. Given the *Desert Thesis*, and given widely accepted norms of interpretative charity—e.g., that (when possible) we interpret others' actions such that they are in general rational and justified—it's plausible that a conventional implicature is created such that, by seeing *A* blame or otherwise hold *B* responsible for *a*-ing, I am (in ordinary circumstances) licensed to infer that *B* is morally responsible for *a*-ing and to regard him as such.<sup>14</sup> In other words, the thought that only morally responsible agents are to be blamed is so deeply embedded into our moral conventions that if we interpret the blamer's actions with a modicum of charity (which would presume that blamers would, in general, blame only if they weren't in violation of the *Desert Thesis*), then we'll have reason to infer that the blamee is morally responsible.

Accordingly, it seems that *if* in blaming another agent, I am (as a result of some ordinary conventions) implicating that he or she is morally responsible for a wrong action and *if* implicature is governed by the standard of reasonable belief, then so too, blame should be governed by the same epistemic standard. I've just argued that we should accept that when we blame others for their actions we conventionally implicated that they are morally responsible for their actions. And earlier, in defense of (3), I argued that implicatures is, in fact, governed by the standard of reasonable belief. Consequently, we can conclude, as AFI concludes, that it is inappropriate (absent special justification) for *A* to blame *B* for *a*-ing if it is not reasonable for *A* to believe that *B* is morally responsible for *a*-ing. ENB is therefore vindicated.

## 6 Objections and Replies I: The Scope of ENB

Of course, you might still worry that this is all too fast. After all, not all instances of blame are expressed. That is, it is possible for blame someone for their wrongdoing privately *in your heart*. For example, a daughter-in-law who resents her mother-in-law for ruining Thanksgiving dinner is plausibly thought to be blaming the mother-in-law even though she doesn't ever mention it.<sup>15</sup> But if it's possible to blame someone without expressing it, as this sort of case suggests, then it's possible to blame without implicating that the target of your blame is morally responsible for performing a wrong action. This suggests that (4) is not always true, and that blame does not, in all cases, create implicature. So ENB cannot be defended on the basis of AFI.

In response, I would first say that paradigmatically, blame is expressed. And this (no doubt) obvious point, which I considered in isolation earlier (in §II.1.), is important because it shows that, at the very least, in a great many of the paradigmatic cases of blame, we *are* subject to ENB. (This alone, I think suffices to show that ENB is not of limited interest, since, for better

<sup>14</sup> Of course, we also know that people regularly blame others even when they are not justified in doing so. Perhaps this gives us a reason to be more cautious in forming the belief that the blamee is morally responsible (or blameworthy) simply on the basis of what's implicated by blame. Similarly, if I know that you in particular are very unreliable, then the fact you have implicated that *p* may not be a reason for me to believe that *p*. It follows, then, that we can have defeaters for the beliefs that we form (or might form) on the basis of others' blame. Though, in such circumstances, we'll have a reason to refrain from blaming that's independent of ENB as I'm applying it here. In ordinary circumstances, however, it seems plausible that credulity is a reasonable response to others' explicit testimony and implicature. Yet even in these circumstances, I want to emphasize, ENB should govern our blame in a way that encourages some resistance to being overly credulous in these contexts.

<sup>15</sup> This example is due to Coleen Macnamara.



or for worse, expressed blame plays such a large role in our social lives.) In other words, even if we weren't subject to ENB in cases of unexpressed blame, our expressed blame would be inappropriate if we do not know that the target of our blame is, in fact, morally responsible. So even if the scope of ENB is limited, it doesn't change the fact that in a wide range of ordinary cases we should refrain from blaming (or otherwise holding morally responsible) if we do not know that the presumptive wrongdoer is, in fact, morally responsible for her action.

Moreover, as a matter of human psychology, it's actually quite difficult to avoid expressing *some* aspect of your blame. An icy stare or a cold tone can express our blaming attitudes even if we're going out of our way to hide them. And of course, the fact that we might be trying to hide our blaming attitudes does not itself cancel what's implicated by the natural expressions of those attitudes. No doubt, in this case, it's the icy stare or cold tone and *not* some linguistic act on the part of the blamer that, in virtue of some features of the proximate social interaction, implicates the target of blame to be morally responsible. But rather than creating a problem for ENB, in such a case, we should simply conclude that even non-linguistic acts—especially non-linguistic acts that are natural expressions of blaming attitudes—are also subject to ENB. Consequently, even if ENB is limited to outward expressions of blame (be they intentional or unintentional), it nevertheless governs a considerable amount of our blaming activity.

But these replies concede a great deal to the objection that I'm not sure needs to be conceded, since it's not at all obvious to me that ENB is limited strictly to cases of expressed blame. Thus second, in response to this challenge, I would argue that there is almost certainly a parallel norm governing instances of unexpressed blame. After all, judgments or beliefs are, in many ways, assertions we make to ourselves. Just as there is something paradoxical about asserting "*p* but I'm not reasonable in my belief that *p*," there is something paradoxical about believing "*p* but I'm not reasonable in my belief that *p*." This suggests, at the very least, that the norms of assertion and belief are likely to be very similar, if not identical.<sup>16</sup> Thus, unexpressed judgments of blame would be subject to the same norms as those judgments that are expressed when we assert or implicate that another agent is to blame.

To further develop this point, we can consider how an agent's judgment that *p* gives her some reason to act on *p* in the same way that her assertion that *p* seems to give others some reason to act on *p*. In this sense, then the relationship between judgments (and the propositional attitudes represented in the emotions like resentment) and assertions is relevantly similar to the relationship between assertions and implicatures. As a result, if we blame another agent, then we license further judgments—even if they stay in our own head—about the presumptive wrongdoer, her character, the quality of her will, etc. And if we are not reasonable in our belief that she is morally responsible, then we will not be warranted in making such judgments. It would, therefore, also be inappropriate for us to act on the basis of these unwarranted judgments for precisely the same reasons that underwrite the impropriety involved in other violations of ENB.

Of course, it's possible to be skeptical of the claim that beliefs and judgments are relevantly similar to assertions in their epistemic norms. Jessica Brown (2012), for example, has recently offered powerful arguments against this claim.<sup>17</sup> In particular, she has claimed that it's very often the case that we do not meet the epistemic standards on asserting that *p* even though it would be appropriate for us to believe that *p* and thereby use it as a premise in our own

<sup>16</sup> For a defense of this claim, see Michael Huemer (2007). N.B., Huemer discusses the paradoxical nature of Moorean assertions of the form "*p* but I do not know that *p*," and so focuses on the connection between assertion, belief, and knowledge. Presumably, the same relation exists between assertion, belief, and reasonable belief.

<sup>17</sup> I am very grateful to an anonymous referee for alerting me to Brown's paper.

practical reasoning. One case she gives for this, TRAIN, is not too dissimilar to *Wrigley Field*.<sup>18</sup> In this case, an agent has consulted the timetable and come to truly believe that the next train is an express train. However, the non-express train, if it were to be the next one, would get her home a mere ten minutes later. Brown rightly, I think, claims that the agent is reasonable in her belief and does not violate any epistemic norms in using it as a premise in her practical reasoning (e.g., by stepping up to the platform and onto the train). But plausibly, Brown also claims, the agent is not in a position to warrantably assert to a stranger that the train is an express train if it is crucial to the stranger that he takes an express train to his destination. It seems, then, that the conditions under which an agent may act on her belief that  $p$  are not isomorphic to the conditions under which she may assert that  $p$ . And this is a significant problem for AFI if it's to provide a basis for ENB, which I've characterized as a norm that governs both expressed and unexpressed blame.

In the case of TRAIN, I do agree with Brown that the agent should not assert that the next train is in fact an express train even though it is appropriate for her to use that as an action-guiding premise in her own practical reasoning. And this certainly puts significant pressure on the idea that beliefs could inherit their norms from the norms of assertion, which are sure to be more restrictive. But so long as it is inappropriate for one to assert that  $p$  whenever it is also inappropriate for one to believe or judge that  $p$ , then Brown's arguments will not touch the weaker point that I want to make.<sup>19</sup> For even supposing that the norms of assertion (and consequently, the norms of implicatures as well) are significantly more restrictive than the norms of belief such that in cases like TRAIN it might be appropriate for agents to believe a proposition without being in a position to assert that proposition, it's nevertheless true that if the agent is not reasonable in her belief that  $p$ , then she can't be warranted in asserting that  $p$ . So any time it would be inappropriate to believe that  $p$ , it would also be inappropriate to assert that  $p$ . In cases of unexpressed blame, then, it looks like the considerations that make it inappropriate to *express* blame will also be in place. AFI explains that because expressions of blame implicate the blamee's moral responsibility (among other things), we should not blame unless we reasonable in believing that the agent is morally responsible. But of course, even if we do not express blame and so implicate the agent's status as morally responsible only to ourselves, it will be inappropriate to believe this of the blamee without good reasons, since even if nothing else is, belief is certainly governed by an epistemic norm. So we shouldn't (absent special justification) blame the agent, even if such blame is unexpressed without it being reasonable to believe that he or she is morally responsible.

Finally, it's important to note that even if you're (quite reasonably) skeptical that *what it is* to blame someone is simply to make a judgment,<sup>20</sup> it is nevertheless plausible to think that unexpressed blame is governed by ENB. This is easy to see in the case of felt but unexpressed blaming emotions like resentment and indignation. For even on theories of the emotions that do not identify emotions with judgments, our emotions nevertheless represent the world in such a way that they are *fitting* or rational only if they are sensitive to evidence as to how things really are, are reliable guides to how things really are, etc. As such, emotions are also subject to epistemic norms,

<sup>18</sup> For more on TRAIN and related cases, see Brown (2012): 140 ff.

<sup>19</sup> This isn't a problem for Brown's argument as I understand it, which is really going after the idea that she calls "Commonality," which is the thesis that there is a single epistemic standard common to assertion and practical reasoning. My argument here will go through as long as a failure to meet the norms of believing that govern practical reasoning is sufficient to ensure that an agent will also fail to meet the norms that govern assertion, and this can be true even if Commonality is false.

<sup>20</sup> For a strong defense of this claim, see R. Jay Wallace (2011).

and in the case of blaming emotions, I see no reason to think that they are subject to a less stringent epistemic norm than any other form of blame, since even if such emotions are wholly unexpressed, they structure the blamer's dispositions towards the blamee in a way that places significant normative burdens on her. Consequently, even though I agree that many instances of blame are unexpressed, this does not tell against the significance of ENB for our practice of blame.

## 7 Objections and Replies II: ENB and Skepticism

Having sorted out the first objection, let us suppose that ENB is true in an effort to evaluate its implications for our blaming practices. Would it simply lead us to modify some of the more problematic features of those practices? Or would it have more radical and implausible implications—implications that might constitute a *reductio ad absurdum* of ENB?

There is some reason to think that ENB—at least at the level of specification we have been operating at so far—would have radically skeptical implications, at least with respect to our current practices of blaming. To see how ENB could lead to a radical form of skepticism, consider the view, *free will and moral responsibility agnosticism*, recently (and as far as I can tell, independently) developed and defended by Stephen Kearns (2015) and Jeremy Byrd (2010). According to Kearns' and Byrd's agnosticism, we should be agnostic about whether anyone ever instantiates the property of being free or morally responsible with respect to their actions.<sup>21</sup> And if, in particular, they're right in thinking that we *should* be agnostic about whether anyone is ever responsible for his or her actions, then it seems that the truth of ENB would entail that we are never in a position to blame others appropriately. Thus, it looks like when ENB is coupled with Kearns' or Byrd's moral responsibility agnosticism—which itself is not an implausible view—the result is a skeptical one.<sup>22</sup>

But why accept some version of moral responsibility agnosticism (be it the one developed by Kearns, or Byrd, or some other version)? Well, consider your favored set of conditions on moral responsibility. In addition to the falsity of causal determinism, incompatibilists typically cite powers like the ability to do otherwise holding fixed the past and the laws or causal sourcehood.<sup>23</sup> Compatibilists, on the other hand often cite rational capacities like reasons-responsiveness or require that actions reflect an agent's "executive control" over her actions—e.g., that an action issues from the agent's "deep self."<sup>24</sup> Given this widespread disagreement about which particular agential capacities ground agents' status as morally responsible, you might think we should adjust our credence of the truth of any one particular theory downward to reflect the plausibility and cogency of any number of other theories.<sup>25</sup> In other words, if the

<sup>21</sup> Strictly speaking, Kearns' thesis is one concerning free will. But since free will (or control) is a necessary condition on instantiating the property of being morally responsible, it would follow that free will agnosticism entails moral responsibility agnosticism.

<sup>22</sup> Indeed, after defending weak free will agnosticism, Kearns notes that one of its consequences might be that we would need to radically alter our blaming practices (Kearns himself focuses on the reactive attitudes, but the point would generalize). And although he doesn't explicitly invoke ENB or some ENB-like principle, it's only in conjunction with an epistemic norm like ENB that free will agnosticism would have significant implications for our blaming practices. Thus, it looks like Kearns doesn't just defend free will agnosticism but also the kind of skepticism that I am considering here.

<sup>23</sup> Cf. Peter van Inwagen (1983) and Derk Pereboom (2001).

<sup>24</sup> Cf. John Martin Fischer and Mark Ravizza (1998), Harry Frankfurt (1971), and Gary Watson (1975).

<sup>25</sup> In other words, you might favor conciliatory replies to the problem of disagreement. For more on such views, see David Christensen (2007). A classic objection to such views can be found in Thomas Kelly (2005).

moral responsibility debates have really reached a “dialectical stalemate” (as a large number of disputants seem to have conceded) between equally coherent and well-formed theories, then you might think we should all be more reticent about wholeheartedly endorsing one theory over any other.<sup>26</sup> Perhaps, it seems, we would do well to be agnostic about which particular agential capacities (if any) are the basis of our status as responsible agents.

Of course, you might note that even if we’re not in a position to form reasonable beliefs about which particular agential capacities ground morally responsible agency, as long as we are reasonable in our belief that we possess *all* of the capacities that experts have taken to be necessary for responsible agency, then we can be reasonable in our belief that others are morally responsible for their actions. And so, even without being in a position to break the dialectical stalemate, we could nevertheless satisfy ENB.

Unfortunately for this attempted reply, it’s not at all clear that we would be reasonable in thinking that human agents possess all the capacities that experts take to be relevant to moral responsibility. Consider, for example, the incompatibilists’ putative requirement that we be able to do otherwise holding fixed the past and the laws. Are we reasonable in our belief that we do, in fact, have the ability to do otherwise holding fixed the past and the laws?<sup>27</sup> At first blush, it can seem that maybe we are. It sure does seem like nothing outside of my own will constrains how I behave. But upon more careful reflection, it can begin to seem that we have very good reasons to worry about the rational status of such beliefs. For after all, suppose that determinism were true and that we therefore lacked the ability to do otherwise holding fixed the past and the laws. It’s nevertheless plausible that even on the assumption that determinism is true, our *experience* of the world would be no different. And because our experience of being able to otherwise—the sense that we sometimes have that at this very moment, we *really* could do A or B—is consistent with the truth of causal determinism, it seems that our best evidence that we satisfy this condition on being morally responsible agents is legitimately in doubt.

Worrisomely, this argument generalizes.<sup>28</sup> So too, our experience of the world will radically underdetermine our status as *the source of our actions*, as *reasons-responsive* (in the appropriate way), and as enjoying *executive control of our actions*. So not only does it look like we have good reasons to be agnostic about whether we are able to do otherwise holding fixed the past and the laws, it looks like we have good reasons to be agnostic about whether we actually satisfy *any* of the putative conditions on being a morally responsible agent. As Gideon Rosen puts it, “given the opacity of the mind—of other minds and even of one’s own mind—it is almost always unreasonable to place confidence in ... any particular positive judgment of responsibility.”<sup>29</sup> And if this is correct, then together with ENB, we could arrive at the very radical conclusion that we *always* have weighty reasons to refrain from blaming others for their actions.

<sup>26</sup> The original characterization of the debate between compatibilists and incompatibilists as being a “dialectical stalemate” is due to John Martin Fischer (1994).

<sup>27</sup> As an anonymous referee has rightly pointed out to me, this over-intellectualizes the issue. What’s really at stake is whether we are reasonable in our belief that we can do otherwise. And issues about the correct analysis of this power aren’t obviously germane to this question.

<sup>28</sup> Gideon Rosen deploys this precise argument in defense of moral responsibility skepticism. According to Rosen, we are responsible for an action only if that action was the result of clear-eyed akrasia. But we are never in a position to determine whether an agent was genuinely akratic, merely ignorant, or weak-willed in some more banal way. Therefore, we are never in a position to be reasonable in our belief that the agent is morally responsible for his or her actions. For more, see Gideon Rosen (2004). What I say below thus serves as a reply to Rosen’s form of moral responsibility skepticism.

<sup>29</sup> Rosen (2004): 308.

But just how good is this argument? I'm inclined to think that it is less damning for our ordinary practices than it might appear at first glance. After all, this argument looks to be a special version of a skeptical argument against other minds. Of course, in this case, our experience isn't shown to radically underdetermine the existence of other agents per se—just the existence of any morally responsible agents. But most of us aren't particularly worried about the problem of other minds. So why should we be worried about this problem? One issue with this rather facile reply is simply that if we are mistaken about the existence of other minds, then we obviously aren't wronging anyone when we fail to treat the things that seem to be, but aren't, agents with due respect. After all, if we are mistaken about the existence of other minds, then there are no other persons who are deserving of such respect. However, if we are mistaken about others' status as morally responsible agents (and we're not even reasonable in believing others to be morally responsible), then, on the truth of ENB (and also, on the truth of the *Desert Thesis*), we wrong others significantly by blaming them.

But this, of course, does not mean that we cannot appeal to any anti-skeptical strategy on behalf of the reasonableness of our belief in others' status as morally responsible agents. Consider the following two cases familiar from debates about whether (and how) epistemic properties like knowledge, justification, and reasonableness depend on context.<sup>30</sup>

#### Low Stakes

Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks [since in the days before direct deposit was ubiquitous, we had to actually go to brick and mortar banks]. It is not important that they do so, as they have no impending bills. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Realizing that it isn't very important that their paychecks are deposited right away, Hannah says, "I'm reasonably sure that the bank will be open tomorrow, since I was there just two weeks ago on Saturday morning. So we can deposit our paychecks tomorrow morning," (Stanley 2005, 3–4).<sup>31</sup>

#### High Stakes

Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since they have an impending bill coming due, and very little in their account, it is very important that they deposit their paychecks by Saturday. Hannah notes that she was at the bank two weeks before on a Saturday morning, and it was open. But, as Sarah points out, banks do change their hours. Hannah says, "I guess you're right. It's not reasonable to conclude that the bank will be open tomorrow and put off depositing the check," (Stanley 2005, 4).

A natural reaction to these two cases is that Hannah is correct in *Low Stakes* when she concludes that her belief that the bank will be open on Saturday is reasonable even though, given the exact same evidence, it seems that she is also correct when she concludes that her

<sup>30</sup> For a contextualist spin on these cases, see Keith DeRose (2009); for a relativist invariantist spin on these cases, see Jason Stanley (2005).

<sup>31</sup> This version of the case is based on Stanley's low stakes case. However, since I am interested in issues concerning reasonable belief rather than knowledge, I have amended Stanley's case to reflect this concern.

belief that the bank will be open is not reasonable in *High Stakes*. Thus it seems that the practical stakes can affect whether or not an agent is reasonable in her belief.

Now, contextualists (e.g., DeRose 2009) and so-called “interest-relative invariantists” (e.g., Stanley 2005) disagree about the precise details as to why it is that Hannah’s belief is a reasonable one in *Low Stakes* even though it is not reasonable in *High Stakes*. But I think we can safely sidestep this intermural debate because the contextualists and interest-relative invariantists agree about what the correct verdict is in each of these cases. That is, both theories seem to independently converge on what we should say in each of these cases, and thus, even without a full theory of *why* the shift in practical stakes affects the reasonableness of Hannah’s belief, it’s enough for our purposes that we see that such shifts do affect the reasonableness of an agent’s beliefs.

If practical stakes can thus affect the reasonableness of an agent’s belief, then this will have important implications for how we think about the reasonableness of our ordinary belief that agents’ regularly satisfy the conditions on being morally responsible agents. For if blame and its typical expressions are (relatively) low stakes affairs, then perhaps we can be reasonable in our belief that others are morally responsible even though our evidence for the proposition that they meet all of the putative conditions on morally responsible agency falls short of what is required in other, higher stakes contexts. That is, just as Hannah can be reasonable in her belief as long as the stakes are sufficiently low, it might be that our belief that agents are morally responsible for their actions is reasonable when blame or other informal sanctions are at stake.

Since my reply to this objection to ENB hinges on this, let me develop this point in a bit more detail. To begin, recall that the moral responsibility agnostic holds that we should refrain from settling on whether any one is ever morally responsible for his or her actions. Why is this? Well, despite the fact that we can see someone do something bad and know that they had a bad motive (people do sometimes tell us these things), we can’t know for sure that the mechanism from which their action issued was appropriately reasons-responsive, that they had the ability to do otherwise, or that they were the source of their action because there are alternative explanations of their behavior—explanations that would be exculpatory—that we cannot rule out on the basis of our current evidence. However, which alternatives must be ruled out in order to be reasonable in a belief plausibly vary according to whether one is in a high or low practical stakes context. We saw this already in the bank cases, since in low stakes, Hannah didn’t need to be able to rule out the alternative that the bank had changed its hours at some point during the two weeks since she had last been there in order to be reasonable in her belief that the bank will be open. Yet in the high stakes case, she did need to be able to rule out this alternative in order to be reasonable. In my view, then, if blame is a low stakes affair, then one need not be able to rule out all of the alternatives in order to be reasonable in one’s belief that an agent doing something bad maliciously is morally responsible. Consequently, even if our evidence for any agent’s being morally responsible is consistent with alternative, but far-fetched explanations of their behavior that are exculpatory, if blame is a low stakes affair, then we don’t need to be able to rule out all such explanations in order to be reasonable in our belief that the agent is morally responsible.

I admit, of course, that the success of this reply relies on my claim that our blaming practices are a low stakes affair. But why should we think this (beyond the fact that it would give us a convenient reply to skeptical worries of the sort expressed above, which is of course, a reason of the wrong kind)? Well, for one thing, blame and other informal sanctions are typically spontaneous and in many cases, these feelings dissipate almost as quickly as they arise. (Think, for example, of how quickly your resentment towards the guy that rudely pushed

in front of you to get the last seat on the train goes away. How often are you still thinking about him two or three minutes after you've gotten off the train?) This means that quite often, when we are apt targets of blame, we are in for (at most) little more than a quick flash of uncomfortable treatment. For another thing, although blame has a characteristic force that no one enjoys experiencing, it is one thing to be blamed and another thing entirely to be punished. And though I do appreciate the genuinely skeptical implications that the truth of ENB together with moral responsibility agnosticism might have for the legitimacy of criminal punishment (as it's currently understood), which involves very harsh treatment, loss of liberty and livelihood, significant loss of reputation and other social goods, and even extreme cases, loss of life, blame by itself does not really threaten wrongdoers in such fundamental ways.

It's quite plausible, then, that the context in which interpersonal blame occurs is a low stakes context. This means that beliefs which have as their propositional content "*A* is morally responsible for her action"—contents that are implicated when we blame—can therefore be reasonable, even though in other, higher stakes contexts, such beliefs would not be reasonable. Accordingly, it's also plausible to conclude that ENB, together with moral responsibility agnosticism, is not a threat to our ordinary practices, although as I've said, it might be a threat to related practices.

## 8 Concluding Remarks

Having articulated, argued for, and defended ENB, I want to conclude by briefly considering one important normative consequence of this principle. In particular, I want to claim that if ENB is true, then it can help to explain why we typically have weighty reason to extend "the benefit of the doubt" to those we presumptively take to be wrongdoers. That is, it can help explain why a certain kind of interpersonal generosity is itself a virtuous character trait.

Because it is often difficult to be reasonable in our beliefs about whether someone is morally responsible for their actions (independently of the very general considerations adduced by moral responsibility agnostics like Kearns and Byrd), the virtuous individual, who takes seriously ENB, is disinclined to blame in a wide range of circumstances in which most of the rest of us are all too happy to lash out. For example, most of us would be quite upset with a close friend if they forgot our birthdays; even if we never mentioned it, we'd probably harbor some resentment towards our friend. But before we get too carried away, we'd do well to remember that actions that are superficially similar can be explained in some cases by forgetfulness, in others by negligence, and in still others, by pure maliciousness. And while we are very plausibly morally responsible and blameworthy for behaviors issuing from negligence or maliciousness, it's not so clear that we deserve blame simply for forgetting. Things do slip our minds from time to time through no fault of our own, and by itself this indicates neither ill will nor callous indifference. As such, *mere* lapses of memory aren't really the proper objects of blame.

To illustrate this last point, consider that the fact that your friend didn't get you anything for your birthday doesn't itself tell us about the underlying mechanism that motivated the seemingly impolitic gesture. Was she sending you a message that she is mad at you, or did she merely have a really stressful day at the office and fell asleep before she remembered to call you? Was she simply unconcerned with celebrating your day, or was she focused, through no fault of her own, on other things and the fact that it was your birthday just wasn't salient (perhaps because, for example, her mother's health took a drastic turn for the worst). In the

scenarios in which your friend knowingly displays ill will or is culpably indifferent, then she is morally responsible, but this isn't so clear in the other cases. And although you know she didn't get you anything or even acknowledge your birthday, you don't know why. Nor do you have good reason to favor one of the above hypotheses over any of the others. At least not without finding out more. Thus, it's not clear from the omission itself that your friend is morally responsible for failing to get you anything for your birthday.

ENB tells us that in these sorts of cases, rather than being presumptively resentful (though perhaps ready to mollify your resentment if you later discover that your friend isn't responsible) you have weighty reasons to instead withhold blame altogether—at least until you have positive reasons to believe that she purposefully decided not to send the birthday card or that she was culpably indifferent. Therefore, ENB provides a basis (of sorts) for the more general moral norm that we not think the worst of others, and it acts as a check on our natural tendency to blame first and then to ask questions later. In other words, ENB entails that in the context of interpersonal blame, we should, as is the case more generally, look before we leap.<sup>32</sup> Perhaps there is something less viscerally gratifying about this than the alternative, but it's certainly a worthy goal.

**Acknowledgments** For help on these ideas and comments on earlier drafts of this paper, I'd like to thank Zac Bachman, John Martin Fischer, Christopher Franklin, Coleen Macnamara, Michael Nelson, and Philip Swenson.

## References

- Brennan, William J. *In re Winship*, 397 U.S. 358 (1970).
- Brown J (2012) Assertion and practical reasoning: common or diverging epistemic standards. *Philos Phenomenol Res* 84(1):23–157
- Byrd J (2010) Agnosticism about moral responsibility. *Can J Philos* 40(3):411–432
- Christensen D (2007) Epistemology of disagreement: the good news. *Philos Rev* 116(2):187–217
- Coates DJ, Tognazzini NA (2013) “The contours of blame,”. In: Coates DJ, Tognazzini NA (eds) *Blame: its nature and norms*. Oxford University Press, New York, pp 3–26
- Coates DJ, Tognazzini NA (2012) The nature and ethics of blame. *Philos Comp* 7(3):197–207
- Cohen GA (2006) “Casting the first stone: Who Can, and Who Can't, condemn the terrorists?”. *R Inst Philos Suppl* 58(113)
- DeRose K (2009) *The Case for contextualism: knowledge, skepticism, and context*, vol 1. Oxford University Press, New York
- Fischer JM (1994) *The metaphysics of free will*. Blackwell, Cambridge
- Fischer JM, Ravizza M (1998) *Responsibility and control: a theory of moral responsibility*. Cambridge University Press, Cambridge
- Frankfurt H (1971) Freedom of the will and the concept of a person. *J Philos* 68(1):5–20
- Frankfurt H (1986) On bullshit. *Raritan* 6:2
- Franklin CE (2013) “Valuing blame”. In: Justin Coates D, Tognazzini NA (eds) *Blame: its nature and norms*. Oxford University Press, New York, pp 207–223
- Grice P (1975) “Logic and conversation,”. In: Peter C, Morgan JL (eds) *Syntax and Semantics vol. 3: Speech acts*. Academic, Waltham, pp 41–58
- Hawthorne J, Stanley J (2008) Knowledge and action. *J Philos* 105(10):571–590

<sup>32</sup> Although this is apparently in conflict with my earlier contention that interpersonal blame is plausibly understood as a low stakes affair, I do not think there is any deep tension between (i) thinking that the epistemic standards in cases of interpersonal blame are fixed by the relatively lower stakes and (ii) also thinking that the truth of ENB would rationalize the virtue of interpersonal generosity. After all, low stakes standards are not *no* stakes standards. This means that while it is relatively easy to meet those standards (and so satisfy ENB) in low stakes cases, we should still be willing to attend to our reasons for believing someone to be blameworthy in a way that is colloquially captured by the injunction to “look before you leap.”



- Herman B (1989) Murder and mayhem: violence and Kantian casuistry. *Monist* 72:411–431
- Huemer, Michael (2007). “Moore’s Paradox and the Norm of Belief,” *Themes from G. E. Moore*, Susana Nuccetelli and Gary Seay, eds. Oxford: Oxford University Press: 142–157
- Kearns S (2015) Free will agnosticism. *Noûs* 49(2):235–252
- Kelly T (2005) The epistemic significance of disagreement. *Oxford Stud Epistemol* 1:167–196
- Lackey J (2007) Norms of assertion. *Noûs* 41(4):594–626
- Pereboom D (2001) *Living without free will*. Cambridge University Press, Cambridge
- Rosen G (2004) Skepticism about moral responsibility. *Philos Perspect* 18:295–313
- Scanlon TM (2008) *Moral Dimensions: permissibility, meaning, blame*. Harvard University Press, Cambridge
- Smith A (2007) On being responsible and holding responsible. *J Ethics* 11:4
- Stanley J (2005) *Knowledge and practical interests*. Oxford University Press, New York
- Todd P (2012) Manipulation and moral standing: an argument for incompatibilism. *Philosophers’ Imprint* 12:7
- van Inwagen P (1983) *An essay on free will*. Oxford University Press, Oxford
- Wallace RJ (2010) Hypocrisy, moral address, and the equal standing of persons. *Philos Public Affairs* 38(4):307–341
- Wallace RJ (2011) “Dispassionate opprobrium: on blame and the reactive sentiments”. In: Jay Wallace R, Rahul K, Samuel F (eds) *Reasons and recognition: essays on the philosophy of T. M. Scanlon*. Oxford University Press, New York, pp 348–372
- Watson G (1975) Free agency. *J Philos* 72(8):205–220
- Watson G (1987). “Responsibility and the Limits of Evil: Variations on a Strawsonian Theme,” *Responsibility, Character, and the Emotions*, ed. F. Schoeman, Cambridge University Press: 256–286
- Williamson T (1996) Knowing and asserting. *Philos Rev* 114(2):489–523