

Epistemic Blame and the New Evil Demon Problem*

Cristina Ballarini

1 Introduction

Philosophers have devoted a lot of attention to the question of what it takes for a belief to be justified. By comparison, they have devoted relatively little attention to the question of what it takes for a belief to be held blamelessly, or to the question of how what some philosophers call “epistemic blamelessness” differs from epistemic justification.¹ A growing number of philosophers have taken this to be a sorry state of affairs. For at least on the face of it, the distinction between beliefs that are justified and beliefs

* This is a pre-copyedit version of an article published in *Philosophical Studies*. Please cite [published version](#).

¹ Philosophers have devoted more attention to the question of what it takes for a belief to be (epistemically) excused—see, for example, Boulton (2017), Brown (2018a: Chapter 4), Cohen and Comesaña (2013), Gerken (2011), Greco (2019), Littlejohn (forthcoming), Schechter (2017), Williamson (forthcoming). That’s because the question of what it takes for a belief to be excused—and not the question of what it takes for a belief to be blameless per se—is of special interest to externalists about justification. For reasons that will become clear in Section 2.3, my focus here will be on the general concept of ‘blameless belief’, which has received somewhat less attention in the literature—though see Brown (2018b, 2020) and Boulton (2020).

that are unjustified but held blamelessly provides proponents of externalist theories of epistemic justification with a principled response to the intuitions which underwrite one of the most serious objections facing externalism: the New Evil Demon Problem.

In this paper, I want to question whether the distinction between epistemic justification and epistemic blamelessness is really such a promising one for externalists to appeal to. In particular, I want to question whether externalists can motivate their appeal to the distinction by analogizing it to a distinction with which we are already familiar: the distinction between moral justification and moral blamelessness.

The paper is divided into three main sections. In Section 2, I introduce the New Evil Demon Problem and the role that the distinction between justification and blamelessness plays in externalist responses to it. These responses—instances of what I call the “unjustified-but-blameless maneuver”—rely, in part, on the assumption that the distinction between justification and blamelessness is one with which we are already familiar. In Sections 3-5, I consider three ways of drawing the distinction familiar from the moral domain: the first in terms of a connection with reactive attitudes, the second in terms of the distinction between wrongness and wronging, and the third in terms of reasons-responsiveness. All three ways of drawing the distinction, I argue, make it difficult to see how an analogous distinction in the epistemic domain could help externalists explain away the intuitions which underwrite the New Evil Demon Problem. In Section 6, I consider two ways externalists might respond. While there may be other ways to motivate the unjustified-but-blameless maneuver, I conclude that it is difficult to do so via an analogy between epistemic and moral evaluation. Motivating the maneuver, then, looks to be a much less straightforward task than its proponents tend to assume.

2 Appeals to Epistemic Blame

2.1 The New Evil Demon Problem

Let's begin with the problem which appeals to epistemic blame are supposed to solve. The New Evil Demon Problem, as it is often called, was first introduced as a problem for Reliabilist theories of epistemic justification (Lehrer and Cohen 1983; Cohen 1984). Simple Reliabilist theories claim that reliability is necessary for justification: you cannot have a justified belief that p unless that belief was produced by a reliable process. The New Evil Demon Problem presents a counterexample to this necessity claim.

We can construct such a counterexample by considering two subjects: call them Ruth and Dupe.² Ruth and Dupe are a lot alike: they have what seem to be exactly the same beliefs, formed on the basis of subjectively indistinguishable experiences, memories, and processes of reasoning. They are disposed to reason in the same way and find the same things intuitive. Given their similarities, we would expect Ruth's and Dupe's beliefs to be equally justified. There is, however, one big difference between the two subjects: Dupe is the unwitting victim of a malicious, all-powerful Cartesian demon. Ruth is not. Thus, while many of Ruth's beliefs (we might suppose) were produced by highly reliable processes, Dupe's corresponding beliefs were produced by highly unreliable ones.³ This possibility presents an apparent counterexample to simple forms of Reliabilism. If these simple forms of Reliabilism are true, then Ruth's and Dupe's beliefs differ in justificatory status. But this result is extremely counterintuitive: Ruth and

² Thanks to Tez Clark for suggesting the name "Dupe".

³ Formulating the problem gets a little tricky if one accepts certain forms of content externalism. I will set these complications aside for the sake of clarity.

Dupe, remember, have what seem to be exactly the same beliefs, formed on the basis of subjectively indistinguishable experiences, memories, and processes of reasoning. Intuitively, Dupe's beliefs should be no less justified than Ruth's.

While the original New Evil Demon Problem was framed as a counterexample to Reliabilism, in recent years the problem has come to play a much larger role in the debate between internalist and externalist theories of epistemic justification. Here, notice that Reliabilism is not the only theory which predicts a difference in the justificatory status of Ruth's and Dupe's beliefs. The New Evil Demon Problem also presents a counterexample to externalist theories which characterize justification in terms of knowledge (Sutton 2005, 2007; Williamson 2000, forthcoming), truth (Littlejohn 2012, forthcoming), and proper-function (Bergmann 2006). More generally, it is a problem for any theory which holds that Ruth and Dupe's beliefs are not equally justified.

To illustrate, consider accounts of justification put forward by proponents of the knowledge-first program. Sutton, to give one example, writes,

“My view is that a subject's belief that p is justified if and only if he knows that p : justification is knowledge.” (Sutton 2007: 7).

Assume that, in the scenario described above, some of Ruth's beliefs constitute knowledge. Dupe's corresponding beliefs are false, and thus do not constitute knowledge. If Sutton's knowledge-first account of justification is correct, then Ruth and Dupe's beliefs must differ in justificatory status as well. Some of Ruth's beliefs—namely, the ones that constitute knowledge—are justified. Dupe's corresponding beliefs do not constitute knowledge, and so are not justified. But again, this result is extremely

counterintuitive: Ruth and Dupe formed their beliefs on the basis of subjectively indistinguishable experiences, memories, and processes of reasoning. Intuitively, Dupe's beliefs should be no less justified than Ruth's.

The New Evil Demon Problem, then, presents an intuitive counterexample to a number of externalist theories of justification. How might proponents of these externalist theories respond?

2.2 The Unjustified-but-Blameless Maneuver

One popular line of response to the New Evil Demon Problem maintains that Ruth's and Dupe's beliefs are not equally justified. Proponents of this response deny that the New Evil Demon Problem presents a counterexample to their favored brand of externalism, and instead hold that the intuitions which underwrite the New Evil Demon Problem are systematically mistaken. Call proponents of this response "Unwavering Externalists".

Unwavering Externalists defend an error-theory about the intuitions which underwrite the New Evil Demon Problem, one which explains why our intuitions about subjects like Ruth and Dupe are systematically mistaken. To construct such an error-theory, Unwavering Externalists need some way of accommodating our positive intuitions about the justificatory status of Dupe's beliefs without identifying these positive intuitions with Dupe's beliefs being outright justified. Here, the strategy many Unwavering Externalists pursue appeals to a distinction between two ways that a subject can be positively normatively evaluated: as justified in believing that *p*, or merely blameless in believing that *p*.⁴ Both Ruth and Dupe's beliefs, the Unwavering Externalist claims, are held blamelessly. But in Dupe's case, we should not confuse a belief's being held blamelessly with its being outright

⁴ See Boulton (2017), Goldman (1988: 59-60), Littlejohn (forthcoming), Srinivasan (2019: 4), and Williamson (forthcoming).

justified, as the intuitions which underwrite the New Evil Demon Problem suggest. Call this way of responding to the New Evil Demon Problem “the unjustified-but-blameless maneuver”.

The unjustified-but-blameless maneuver appeals to the idea that we can evaluate a belief in terms of how “blameworthy” a subject is for holding it. This immediately raises two questions. First: How should we understand the notion of blameworthy belief, as it is employed by proponents of the unjustified-but-blameless maneuver? Second: Why think that subjects like Dupe are merely blameless, as opposed to outright justified, in believing as they do?

Starting with the first question: The kind of blameworthiness relevant to the unjustified-but-blameless maneuver is supposed to be associated with a kind of epistemic evaluation. When the Unwavering Externalist says that Dupe is blameless for believing as she does, this is supposed to be an assessment of her epistemic status, rather than her moral status. To understand the distinction, consider the following three cases.

RACIST: At a post-conference dinner, Albert, a white man, asks his fellow conference-goer Bill, a Black man, to refill his drink. The request reflects a certain belief Albert holds: that, rather than a fellow academic, Bill is instead a member of the wait staff.

CLAIRVOYANT: On the basis of her completely reliable clairvoyant power, Maud forms the belief that the President is in New York. She maintains this belief in the face of excellent evidence against the very possibility of clairvoyance.

CONSPIRACY THEORIST: Carl believes that the Earth is flat and that widespread acceptance of its sphericity is part of a massive

global conspiracy. When presented with overwhelming evidence to the contrary, he either disregards this evidence or attempts (rather convolutedly) to explain that evidence away.

The three cases above all involve subjects who are, in some way or another, criticizable for believing as they do. Here, many philosophers have found it natural to describe what is criticizable about such subjects in terms of their being blameworthy for believing as they do.⁵ We can distinguish, however, between two reasons why subjects might be considered blameworthy in such cases.

Cases of apparently blameworthy belief like RACIST involve a kind of blame which is clearly moral in flavor: if Bill's blaming Albert for his belief is appropriate, it is appropriate because of a certain kind of moral failing on Albert's part.⁶ In what follows, I'll refer to the kind of blame at issue in RACIST as 'moral blame'. Cases like CLAIRVOYANT and CONSPIRACY THEORIST may also involve moral blameworthiness, but they need not. Even stipulating that Maud and Carl's beliefs are in no way morally problematic, many philosophers have felt that Maud and Carl are still, in some other sense, blameworthy for holding them.⁷ If blame in these cases is appropriate, these philosophers have thought, it is appropriate because of a certain kind of *epistemic* failing on the subject's part. According to these philosophers, Maud and Carl are blameworthy for manifesting

⁵ The first case is discussed by Basu (2019a). The second case, from Bonjour (1980), is discussed in terms of blame by Brown (2018b) and Hawthorne and Srinivasan (2013).

⁶ Though no one would be particularly surprised if Bill did blame Albert for his hurtful belief, the question of whether Bill would be justified in blaming Albert is a controversial one among philosophers. For arguments in favor of the thesis that we can be held morally accountable (e.g. blamed) for our beliefs, see Basu (2018, 2019a, 2019b), Hieronymi (2008), McHugh (2013, 2014), and Smith (2005). For arguments against, see Levy (2007) and Setiya (2013).

⁷ See Brown (2018b) for discussion.

dogmatism, a kind of poor epistemic character. In what follows, I'll refer to the kind of blameworthiness at issue in CLAIRVOYANT and CONSPIRACY THEORIST as 'epistemic blame'.⁸

For the Unwavering Externalist, the notion of blameworthy belief relevant to the New Evil Demon Problem is to be spelled out in terms of epistemic blame. Subjects like Dupe, the Unwavering Externalist thinks, are epistemically blameless for believing as they do in the same way that subjects like Maud and Carl are epistemically blameworthy for believing as they do. This is not only explicit in how proponents of the unjustified-but-blameless maneuver frame their appeals to epistemic blame,⁹ but is also clear if we consider what the intuitions which Unwavering Externalists are trying to explain away seem to track. On the face of it, our positive intuitions about Dupe's beliefs are intuitions about Dupe's epistemic status. They do

⁸ Before moving on, I should note that many of the cases discussed in the literature on blameworthy doxastic attitudes involve subjects who are intuitively blameworthy for failures to believe (or know) some relevant proposition. In the moral case, consider an intuitive example from Smith (2005): it seems appropriate for my close friend to morally blame me for failing to remember her birthday, a date which I know is very important to her. In the epistemic case, Weatherson (2008) writes that there is something epistemically blameworthy about a subject who, after looking in the fridge, fails to believe that the house needs more orange juice because he fails to check whether the carton he sees—which he knows his roommate tends to put back in the fridge, even when empty—is also empty on this particular occasion. (Goldberg [2017] discusses similar examples but defends the view that the fact that subjects should have known some relevant proposition in such cases makes them unjustified in their beliefs.) The idea that subjects might be blameworthy not only for holding morally or epistemically problematic beliefs, but also for failing to believe certain morally or epistemically significant propositions, brings up a number of interesting issues. These issues, I think, are irrelevant for my purposes here. In what follows, then, I will focus on instances of purportedly blameworthy/blameless belief, though I suspect my arguments will generalize to other instances of purportedly blameworthy doxastic conduct.

⁹ See Littlejohn (forthcoming), Hawthorne and Srinivasan (2013), and Williamson (forthcoming).

not seem to depend on any further assumptions about the moral details of Dupe's situation.

This brings us to our second question: Assuming that there is a notion of blameworthy belief relevant to the New Evil Demon Problem, why think that Dupe's beliefs are *merely* epistemically blameless, as opposed to outright epistemically justified? A popular way of motivating the unjustified-but-blameless maneuver appeals to the notion of excuse (Boult 2017; Littlejohn forthcoming; Williamson forthcoming). Our positive intuitions about subjects like Dupe, the Unwavering Externalist claims, are sensitive to the fact that she has an excellent excuse for holding the beliefs that she does: we are sensitive to the fact that Dupe has been unwittingly deceived, is in unusual circumstances, etc. But while having an excuse is sufficient for being blameless, it is not sufficient for being justified. In general, the Unwavering Externalist claims, subjects only need an excuse when they have failed to meet the standards required for justification.¹⁰

¹⁰ Externalists, then, are interested in excuse, not just blameless belief per se. It's worth pausing to think about why this is the case. The category of blameless belief is a quite diverse one: plausibly, it includes both false beliefs like Dupe's—beliefs impeccably formed on the basis of misleading evidence—but also false beliefs formed by subjects suffering from temporary insanity or incapacitation. Externalists don't want to say that Dupe's beliefs are blameless in the same way as the beliefs of the insane and incapacitated. For, as many authors have already noted, we're supposed to have stipulated that Dupe's beliefs were formed impeccably on the basis of her misleading evidence—the same cannot be said of beliefs formed by those whose belief-forming capacities have been compromised. So it would be misleading to lump Dupe's false beliefs in with the false beliefs of the insane and incapacitated (see Cohen and Comesaña [2013], Madison [2018: 4559], Schechter [2017], and Pryor [2001: 117] for versions of this criticism). Here's where the externalist's limited focus comes into play: Plausibly, beliefs formed by incapacitated or insane subjects are exempted, but not excused, from blame (see Littlejohn [forthcoming] for further discussion of this distinction). So externalists will object to the comparison between these beliefs and the ones held by subjects like Dupe. (I thank an anonymous reviewer for pressing me on this point.)

Unwavering Externalists often motivate this picture by locating it within a general normative framework expressed in terms of norm compliance and violation (Littlejohn 2012, Williamson forthcoming). In general, these Unwavering Externalists claim, subjects are justified if and only if they comply with the norms relevant to their situation. A subject is unjustified if and only if they violate these norms. In the epistemic case, externalists claim, the norms for belief reference conditions which Dupe does not meet. So Dupe's beliefs are unjustified. Even so, Unwavering Externalists point out, there are better and worse ways for a subject's belief to fall short of justification. The better way involves the subject's being blameless for violating the relevant epistemic norms, because that subject has an excuse for doing so. In this case, even though the subject did not comply with the primary norm at issue, she might have complied with some derivative norm to which it gives rise (e.g. "Do as the person who tends to comply with the primary norm would!"). The worse way involves the subject's being epistemically blameworthy for violating the relevant norms. In this case, the subject complies neither with the primary norm nor the derivative norm(s) to which it gives rise.¹¹

2.3 Assessing the Unjustified-but-Blameless Maneuver

¹¹ To be more careful: The externalist will say that there's some sense in which failing to meet a normative standard with an excuse or exemption is better than failing to meet that normative standard without any excuse or exemption. There's of course another sense in which it's not obviously better: Consider the case where one is exempted from blame because one lacks the capacities to be held responsible. In this case, one has (1) failed to meet a normative standard and (2) lacks the capacities to be held responsible. In some sense, it would be "better" if only (1) were true. For further discussion of the "ladder of defences" in the legal context, see Gardner (2007: Section 3). (Thanks to an anonymous reviewer for pressing me to be more careful here.)

What should we think about the unjustified-but-blameless maneuver? Most critical discussion of the maneuver has focused specifically on the distinction between justification and excuse. Here, a number of philosophers have questioned the claim that Dupe’s beliefs are merely excusable, rather than outright justified (and thus in no need of an excuse). I think that this question is an interesting and important one, especially given that the most popular way of explaining why subjects like Dupe merely blameless in believing as they do appeals to the notion of excuse.¹² In what follows, however, I want to take a step back from talk of justification and excuse to focus on an assumption operating in the background such debates: the idea that Dupe might be “epistemically blameless” for believing as she does.

One reason why the unjustified-but-blameless maneuver seems well-motivated, in general, is that the distinction between justification and blamelessness strikes us as a very familiar one. We are quite comfortable talking about blame and blamelessness in the moral domain. So it seems kosher, absent any special reason to think morality and epistemology would not run in parallel here, to talk about blame and blamelessness in the epistemic domain. The thought that the unjustified-but-blameless maneuver appeals to a very familiar, though too often overlooked, normative distinction is apparent in how Unwavering Externalists often frame the

¹² See Boulton (2017), Cohen and Comesaña (2013), Gerken (2011), Greco (2019) and Schechter (2017) for relevant discussion. One thing that complicates the debate, here, is that Unwavering Externalists have at their disposal what Greco calls the “Anna Karenina Response”: “[j]ustified beliefs are all alike...but each merely excused belief is merely excused in its own way” (Greco 2019: 4). So a natural strategy for replying to the unjustified-but-blameless maneuver—pointing to systematic differences between Dupe’s case and other cases of merely excusable epistemic activity—is not available. For further discussion, see Littlejohn’s (forthcoming) discussion of “contrast arguments”. If the Anna Karenina Response is a good one—and I think that it is—opponents of the unjustified-but-blameless maneuver would do well to appeal to independent considerations (e.g. considerations about blame). That is my strategy here.

maneuver.¹³ Williamson, for instance, introduces his appeal to the maneuver as follows:

Epistemologists routinely use normative terms such as ‘justified’, ‘rational’, and ‘reasonable’, but too often neglect vitally relevant, though subtle, normative distinctions... This essay steps back from the epistemological issues to make some of those general normative distinctions, then returns with them to epistemology. (Williamson, forthcoming)

It is worth asking, then, whether the distinction between justification and blamelessness with which we are already familiar can really be so neatly imported over from discussions in moral philosophy; if it cannot, then it is not clear that the unjustified-but-blameless maneuver is well-motivated. To do so, we should start by considering the role that blame plays in moral theorizing. What are we saying about a subject when we say that her action or belief is morally blameworthy or blameless, as opposed to morally justified or unjustified? In what follows, I’ll consider three ways of drawing the distinction between justification and blamelessness in the moral domain. In each case, I’ll argue, it is difficult to see how an analogous distinction in the epistemic domain could help the externalist explain away the intuitions which underwrite the New Evil Demon Problem.¹⁴

¹³ That blame in the moral and epistemic cases should be treated in a unified fashion is asserted both by proponents of strongly externalist theories—see Williamson (forthcoming) and Littlejohn (forthcoming)—and by those more critical of such theories—see Brown (2018b).

¹⁴ Before moving on, I encourage the reader to note that evaluation in terms of blameworthiness is typically taken to presuppose that the subject being evaluated can be held responsible for their action or attitude. Some philosophers think that it makes no sense to consider subjects responsible or accountable for their beliefs, since we have no voluntary control over our beliefs and since voluntary control is required for responsibility or

Before presenting these arguments, let me be clear about what I am trying to achieve with them. I do *not* mean to argue, in what follows, that there is no systematic way of drawing the distinction between epistemic justification and epistemic blameworthiness that would deliver the result that subjects like Dupe are epistemically unjustified but blameless. In fact, I think that there might be a number of ways to fill in the idea that Dupe is epistemically unjustified but X, where X is some other positive normative status that falls short of justification. What I want to challenge is the idea that we should understand X as *blamelessness*, and so the idea that the externalist's error-theory can be partly motivated via the moral analogy. If I'm right, then the Unwavering Externalist needs to motivate this error-theory on independent grounds.

Relatedly: In recent years, epistemologists have been increasingly interested in epistemic blame. I do *not* mean to argue that these epistemologists are wasting their time, or that there is no notion of epistemic blame useful for epistemological theorizing. What I want to emphasize, rather, is an apparent disconnect in the literature on epistemic blame. One reason for the increasing interest in epistemic blame has been an increasing interest in the social dimensions of epistemic normativity. Another reason has been an increasing popularity of strongly externalist theories of epistemic justification, many of which make use of something like the unjustified-but-blameless maneuver. While the accounts of epistemic blame developed by social epistemologists may be independently interesting for

accountability (see Levy [2007] and Setiya [2013]). Of course, many philosophers deny this, holding that subjects can be responsible or accountable for things they lack voluntary control over (see Hieronymi [2008], McHugh [2013, 2014], and Smith [2005]). And I myself will proceed with this assumption. But we should note that the assumption is not uncontroversial. Thus, proponents of the unjustified-but-blameless maneuver who do take voluntary control to be required for responsibility will need to rethink their appeal to epistemic blame at a much earlier step in the dialectic than I will be supposing.

epistemology as a discipline, such accounts do not (I want to suggest) fit nicely with the accounts of epistemic blame that Unwavering Externalists need to motivate the unjustified-but-blameless maneuver. With these clarificatory remarks in mind, let's now consider some roles that the distinction between justification and blamelessness plays in moral theorizing.

3 Blame and Reactive Attitudes

In this section, I'll consider one role that the distinction between justification and blamelessness plays in moral theorizing. I'll argue that Unwavering Externalists cannot motivate the unjustified-but-blameless maneuver by drawing an analogous distinction in the epistemic domain. In Sections 4 and 5, I'll consider two other roles that the distinction between justification and blamelessness plays in moral theorizing. In those cases, too, I'll argue that Unwavering Externalists cannot motivate the unjustified-but-blameless maneuver by drawing an analogous distinction in the epistemic domain.

3.1 Moral Blame and Reactive Attitudes

In epistemology, philosophers have devoted much more attention to questions about epistemic justification than they have to questions about epistemic blame and blameworthiness. In moral philosophy, the focus has been quite different. While philosophers have written a lot about moral

blame and blameworthiness, justification is not a term one hears about a lot in ethics, outside of the legal context.¹⁵

Still, there is a kind of evaluation which we might call “evaluation in terms of moral justification”, which bears a close connection to the familiar act of “justifying oneself”. In the law and in everyday life, we often try to justify our actions and attitudes, whether on moral or legal grounds. When we do so, the idea is to show that, given the details of our situation, our action or attitude was morally or legally right, rather than morally or legally wrong. Classic justificatory defenses include self-defense, necessity, and authority. To give an example from Greco (2019: 7): If I smash your car window because doing so is the only way to save the infant trapped inside from overheating, that action is justified by necessity. Though in the legal case, giving such a justificatory defense involves claiming that something which would usually be illegal was, given the total details of the situation, legal, we can imagine analogous sorts of defenses made on purely moral grounds. In such cases, claiming that the subject’s action was justified by necessity involves claiming that, though the subject’s action was an instance of the sort of action which is generally considered morally wrong, given the total details of the situation, it was in fact the morally right thing to do. So

¹⁵ There are, of course, a few prominent exceptions. T. M. Scanlon’s contractualism, in its original formulation, equates wrongness with unjustifiability (Scanlon 1998). And (as an anonymous reviewer helpfully points out), there’s this passage from G. E. Moore (1903: §60):

The only possible reason that can justify any action is that by it the greatest possible amount of what is good absolutely should be realized. And if anyone says that the attainment of his own happiness justifies his actions, he must mean that this is the greatest possible amount of Universal Good which he can realise.

So the term ‘justification’ does show up occasionally in moral philosophy. When it does, the justified is often treated as a rough equivalent of the right, as the passage above suggests. This is all compatible with how I’ll go on to characterize moral justification.

a subject's action or attitude is morally justified, we might say, if that subject's action or attitude is morally justifiable. On this account, moral justification is associated with (though not obviously the same thing as) an action or attitude's being morally right.

In Sections 3-5, I'll assume that something like this account of justification in the moral domain is roughly correct. What, then, does the distinction between justification and blamelessness amount to? One way to draw this distinction is by appeal to what are called "reactive attitudes". Moral justification, as I claimed above, is associated with an action or attitude being morally right. Blamelessness, on the other hand, is associated with the fact that it would be inappropriate (not "fitting") to adopt a certain set of negative attitudes and emotional reactions toward the subject being evaluated.

This way of drawing the distinction goes with a familiar picture of what it means to be morally blameworthy or blameless. To get a feel for this picture, it will help to note the distinction between the property of blameworthiness, the activity of blaming someone, and judgments about someone's being blameworthy. The property of blameworthiness, on this natural picture, is associated with the appropriateness (or "fittingness") of certain negative attitudes and emotional reactions. To illustrate: The fact that I intentionally stepped on your toe makes fitting certain kinds of reactions on your part. These reactions might include some combination of "hot reactive attitudes": negative emotional reactions such as your feeling angry or resentful toward me. They may also include some combination of "cold reactive attitudes": conative adjustments such as the withdrawal of trust, a revision of your desire to spend time with me next weekend, or a decreased willingness to assist me in my projects.¹⁶ In such a case, you might

¹⁶ There is some debate over just which reactive attitudes are required for a subject to count as blaming another. Emotional theories of blame stress the hot reactive attitudes

actually blame me, adopting some combination of these reactive attitudes. Or you might merely judge me blameworthy (that is, judge that these reactive attitudes are fitting) without actually blaming me—perhaps because you recognize that I already feel guilty for my callousness and am making efforts to change (Smith 2007: 482).

The idea that blamelessness is associated with the inappropriateness of adopting certain negative attitudes and emotional reactions toward the subject being evaluated is implicit in much of the literature on moral blame and blamelessness.¹⁷ To this point, many philosophers characterize blame in terms of reactive attitudes directly. To take a few examples:

In one way, to blame (morally) is to attribute something to a (moral) fault in the agent; therefore, to call conduct shoddy *is* to blame the agent. But judgments of blameworthiness are also thought to involve the idea that agents deserve adverse treatment or “negative attitudes” in response to their faulty conduct. (Watson 1996: 230-231)

My ordinary use of the term associates blame with a certain kind of negative emotional attitude toward the object of blame—resentment, indignation, and guilt, as well as righteous anger, fall within the family of these attitudes... (Wolf 2011: 335)

(see Strawson [1962], Wallace [1994], and Wolf [2011]). Conative theories of blame stress the colder reactive attitudes (see Sher [2005] and Scanlon [2008, 2013a]).

¹⁷ Connecting this with the work on excuses, there are a few different reasons why negative reactive attitudes might not be fitting: a subject’s action may be outright justified, unjustified but excused, or the subject may be exempted from blame (despite the fact that her action was both unjustified and unexcused).

To *blame* a person is to judge that person to be blameworthy and, as a consequence, to modify one's understanding of one's relationship with that person (that is, to alter or withhold intentions and expectations that that relationship would normally involve) in the particular ways that that judgment of blameworthiness makes appropriate, given one's relation with the person and the significance for one of what that person has done. (Scanlon 2013a: 89)

Of course, not all philosophers characterize blame and blameworthiness in terms of reactive attitudes directly. Philosophers who endorse "cognitive" accounts of blame, for instance, characterize blame as a kind of negative assessment of a subject's action or attitude and the character it reflects.¹⁸ Still, advocates of cognitive accounts think that blame and blameworthiness bear a close conceptual connection to the reactive attitudes: blameworthiness is associated with (though not obviously the same thing as) the fittingness of reactive attitudes.¹⁹

Contrast this with the case of justification. An action or attitude can have the property of being morally justified or unjustified. And we can make judgments about whether a given action or attitude has either of these properties. But whether or not such judgments are apt will not necessarily involve the fittingness of any particular set of reactive attitudes. Moral justification is a matter of whether one's action or attitude was all-things-considered morally right, or whether it conformed with the relevant moral standards or norms. That is why, in judging that my intentionally stepping

¹⁸ Early versions of such accounts characterize identify blame with a negative judgment about a subject's "moral ledger" (see Smart [1961] and Zimmerman [1988]). More contemporary versions identify blame with the judgment that a subject's action manifested insufficient responsiveness to moral reasons (see Arpaly [2006] and Hieronymi [2004]). I will come back to such accounts in Section 5.

¹⁹ See, for instance, Arpaly (2006: 25-28).

on your toe was unjustified, we appeal to facts about what is morally right or wrong, or the moral norms relevant to the situation (i.e. it's morally wrong to intentionally cause a person pain). These facts are associated with my action being unjustified. While it may be fitting to adopt certain reactive attitudes toward me for unjustifiably stepping on your toe, this is far from generally the case. Young children, people suffering from temporary insanity, and individuals who are otherwise mentally incapacitated can perform actions which are morally unjustified—like intentionally stepping on someone's toe. But it would not be fitting, many people have thought, to adopt any blame-related reactive attitudes toward such subjects for their transgressions.

3.2 Epistemic Blame and Reactive Attitudes

Let's see if this way of drawing the distinction between justification and blamelessness can help motivate the unjustified-but-blameless maneuver. Unwavering Externalists, remember, want to understand the distinction between epistemic justification and blamelessness as something analogous to the distinction between moral justification and blamelessness. If they characterize this distinction in terms of reactive attitudes, we get the following result: Epistemic justification is associated with a belief's being "the right one" to hold, from the perspective of some set of epistemic standards or norms. Epistemic blamelessness, in contrast, is associated with the fact that it would not be fitting to adopt some relevant set of reactive attitudes toward that subject for believing as she does.

I think that it's unclear how Unwavering Externalists might characterize the reactive attitudes relevant to epistemic blame in a way which motivates the unjustified-but-blameless maneuver. To appreciate the

problem, we should keep in mind the cases which the unjustified-but-blameless maneuver needs to cover. Unwavering Externalists claim that:

- (1) Our positive intuitions about the epistemic status of Dupe's beliefs are explained by the fact that she is epistemically blameless for holding the beliefs that she does.
- (2) The negative correlate of these intuitions is at work in the case of BonJour's Maud and conspiracy theorist Carl. Unlike Dupe, Maud and Carl are epistemically blameworthy for holding the beliefs that they do.

If we are to understand blameworthiness in terms of reactive attitudes, then (1) and (2) add up to the claim that our positive intuitions about the epistemic status of Dupe's beliefs are explained by the fact that it would not be fitting to adopt some set of reactive attitudes toward her—attitudes it would be fitting to adopt toward subject's like Maud and Carl.²⁰

What might these reactive attitudes be, in the epistemic case? We can start by considering the kinds of hot reactive attitudes we standardly associate with a subject's being morally blameworthy: attitudes like anger and resentment. Understanding blame in terms of these hot reactive attitudes seems to get Dupe's case right: It does not seem fitting to get mad

²⁰ Note that (1) should be understood in a very specific way. The externalist does want to say that Dupe is epistemically blameless while Maud and Carl are epistemically blameworthy. But for the externalist, the considerations that can render a subject blameless are a "motley bunch", as Littlejohn (forthcoming) puts it (see also my discussion in Footnote 9). Some considerations—like non-culpable ignorance—excuse a subject from blame. Other considerations—like infancy or insanity—exempt a subject from blame. Going forward, then, it's important to note that what most externalists writing about these issues have in mind with (1) is not the kind of blamelessness associated with infancy, insanity, and exemption—what they have in mind is the kind of blamelessness associated with ignorance and excuse. I thank an anonymous reviewer for pressing this point.

at or resent Dupe for holding the beliefs that she does. For even if these beliefs fall short in a way that Dupe's counterpart Ruth's beliefs do not, they only fall short because Dupe has been systematically misled. So understanding blameworthiness in terms of hot reactive attitudes gets us (1): that Dupe is epistemically blameless. Understanding blameworthiness in terms of reactive attitudes like blame and resentment, however, seems to get the cases of Maud and Carl wrong. For while it seems appropriate to get angry at or resent subjects on the basis of their moral failings, it seems inappropriate to get angry at or resent them *purely on the basis of their epistemic failings*. While we might in fact get frustrated with subjects like Maud and Carl for other reasons, it is not obvious that anger or resentment is a fitting response to their dogmatism per se. Understanding blameworthiness in terms of hot reactive attitudes, then, fails to secure claim (2): that Maud and Carl are epistemically blameworthy.

This last point is a bit of a delicate one. In many everyday cases, of course, we do get mad at (and perhaps even resent) dogmatists like Maud and Carl. The question we should ask ourselves, though, is whether such anger and resentment seem fitting because of Maud and Carl's epistemic failings or because of the moral failings with which these epistemic failings are often associated. I think that the more we stipulate that Maud and Carl's dogmatism is not morally problematic—they have not unjustly dismissed anybody's testimony, they lack the capacity to act on their beliefs, and so on—the less anger and resentment seem like fitting responses. There may be some special contexts in which it is appropriate to feel anger and resentment toward people purely on the basis of their epistemic failings—for instance, in professional contexts where others depend on one's being knowledgeable about a certain range of topics (Goldberg 2017, 2018). But it is at best controversial whether the fittingness of such responses applies outside of such contexts (Hazlett 2018; Brown 2020). To this point: Consider

the wide range of beliefs that Dupe is supposed to strike us as blameless for holding. These include beliefs arrived at through effortful processes of reasoning, but also beliefs arrived at via perception and memory. Do our positive intuitions about the epistemic status of these beliefs really pick up on the fact that it would be inappropriate to get angry at or resent Dupe for holding them? Here, I think, the answer is clearly no.

How else might Unwavering Externalists characterize the reactive attitudes relevant to epistemic blame? While purely epistemic failings may not warrant emotions like anger or resentment, they do seem to warrant certain changes in one's general orientation toward the subject being evaluated. If I learn that you tend to base your beliefs on insufficient evidence, it seems appropriate for me to reduce my confidence in your testimony, my willingness to treat you as a partner in shared inquiry, and so on.²¹ More realistically, then, Unwavering Externalists might characterize the reactive attitudes relevant to epistemic blame in terms of what I called the cold reactive attitudes.

This account of the reactive attitudes seems to get cases like Maud and Carl's right. Plausibly, the kind of dogmatism Maud and Carl's beliefs manifest makes it appropriate to reduce one's confidence in their testimony, one's willingness to treat them as partners in shared inquiry, and so on. So this way of characterizing the reactive attitudes relevant to epistemic blame secures claim (2): that Maud and Carl are epistemically blameworthy for believing as they do. But the Unwavering Externalist, remember, also wants to claim that Dupe, unlike Maud and Carl, is epistemically blameless for believing as she does. If epistemic blame involves reducing one's confidence in a subject's testimony, one's willingness to treat them as a partner in shared inquiry, and so on, then it would seem that Dupe is epistemically

²¹ See Boulton (2020) and Kauppinen (2018). Note that Kauppinen does not attempt to associate these reactive attitudes with epistemic blame.

blameworthy too, just like Maud and Carl. To this point, consider the wide variety of reasons why one might adopt these epistemically tinged reactive attitudes. One reason is that the subject in question has poor epistemic dispositions, as is the case for subjects like Maud and Carl. But another reason is that the subject in question is poorly situated with respect to the facts one is interested in learning. If I want to know whether someone's remarks toward me were sexist, it seems perfectly fitting to have low confidence in, or decide not to consult, my elderly neighbor—a man who has never experienced sexism himself and who, for various reasons, has never learned to reliably identify instances of sexism. Dupe's case is a rather extreme version of this latter case: Dupe is poorly situated with respect to virtually all facts one might be interested in learning. So it seems just as appropriate to adopt these colder reactive attitudes toward subjects like Dupe—who form systematically false beliefs because of their unfortunate epistemic situation—as it does to adopt them toward subjects like Maud and Carl—who form systematically false beliefs because of their bad epistemic characters. Understanding blameworthiness in terms of these cold reactive attitudes, then, fails to secure claim (1): that Dupe is epistemically blameless.²²

Now, the Unwavering Externalist might object to my description of the cases above. I'm right, they might say, that blameworthiness is associated with the fittingness of certain reactive attitudes. But it is not obvious, in these cases, that it's appropriate to adopt the relevant reactive attitudes because they are *fitting responses to blameworthiness*. In the moral case, it can be appropriate to adopt blame-related reactive attitudes for reasons other than the fact that these are fitting responses to a subject's

²² I'm not the first to notice that Dupe, Maud, and Carl can be difficult to distinguish given things externalists say about justification and blamelessness—see, for example, the discussion in Simion, Kelp, and Ghijsen (2016).

blameworthy conduct. Returning to the example above, if I intentionally step on your toe because I was somehow tricked into thinking that is what you wanted, I have an excuse—and am therefore morally blameless—for doing so. You might have pragmatic reasons to get mad at, resent, or decrease the trust you place in me (perhaps doing so would be an effective mode of deterrence). But these responses would only be appropriate in a pragmatic sense—they would not be fitting responses to my action's being morally blameworthy. In the case above, it's not clear that it would be fitting (rather than just pragmatically appropriate) to reduce one's confidence in Dupe's testimony, one's willingness to treat her as a partner in shared inquiry, and so on. If that's right, then the externalist can still claim that Dupe is unjustified but blameless.

Let's try to fill out this line of thought. How can we tell whether reactive attitudes are appropriate to adopt because they are fitting responses to blameworthiness, as opposed to their being merely pragmatically beneficial? To do so, we must attend to facts about the subject whose faulty conduct is being evaluated. In the moral case, the externalist might claim, reactive attitudes do not count as responses to blameworthiness if they are directed at subjects whose actions or attitudes were *merely the result of bad circumstances*. In the example above, blame-related reactive attitudes would not be fitting responses to my intentionally stepping on your toe because my doing so was merely the result of my being tricked into doing it. We can apply this idea, the externalist might claim, to the epistemic case. While there are many reasons why we may decrease our confidence in a subject's testimony, our willingness to treat her as a partner in shared inquiry, and so on, these responses should not be associated with epistemic blameworthiness if the subject's faulty beliefs merely resulted from her poor epistemic circumstances. Since Dupe's beliefs are only faulty because she

has been unknowingly deceived, we can maintain that she is epistemically blameless for believing as she does.

I think that there's something to this line of thought. But it's ultimately not a promising one for the externalist to pursue. To see why, first consider the supposed analogy with the moral case.²³ The claim that subjects are morally blameless for actions or attitudes which were merely the result of bad circumstances is a bit vague on its own: whether this claim is true will depend on the principle that's supposed to be behind it. If the idea is that subjects can't be blamed for actions or attitudes that are merely result of bad circumstances *because these circumstances are ones which are not in the subject's control*, then that claim is very controversial. A subject might have a morally problematic character or disposition because of factors outside of her control. Many philosophers would maintain, though, that such subjects can still be blameworthy for actions or attitudes which manifest this bad character.²⁴

More plausibly, then, the externalist might claim that subjects can't be morally blamed for actions or attitudes that are merely the result of bad circumstances *because such actions and attitudes are not manifestations of the subject's bad character*. But this version of the response won't help the externalist much either. To see why, consider how this principle about character is supposed to carry over to the epistemic case. Suppose the Unwavering Externalist claims that Dupe is epistemically blameless because her faulty beliefs are not manifestations of a bad epistemic character. What does it mean to manifest a bad epistemic character? Plausibly, beliefs manifest a bad epistemic character if they are produced by a process which

²³ Thanks to Evan Behrle for discussion.

²⁴ For arguments to this effect, see Scanlon (2008: 194-198; 2013b: 108-109; 2015: 97-105). See also Footnote 11.

tends to produce bad epistemic results—for example, a high ratio of falsehoods to truths. So the externalist’s proposal now looks like this:

- (1) Dupe is epistemically blameless because her faulty beliefs did not result from a process which reliably produces bad epistemic results.
- (2) Maud and Carl are epistemically blameworthy because their faulty beliefs did result from a process which reliably produces bad epistemic results.

At this point, the Unwavering Externalist will need to specify how she is evaluating the belief-forming processes in question. Are the belief-formation processes associated with bad epistemic character ones which tend to produce bad results in all possible worlds? In “normal worlds”? Or in the subject’s world? Consider these options in turn.

The first option seems implausible on its face. Suppose the belief-forming processes associated with bad epistemic character are ones which tend to produce bad results in all possible worlds. Not only Dupe but Ruth, too, would be blameworthy on this reading: they both form beliefs using belief-forming processes which would tend to produce bad epistemic results in the demon world. This is obviously not the result that the Unwavering Externalist wants.

The Unwavering Externalist, then, will probably want to relativize the kind of character assessment associated with blameworthiness. Suppose she relativizes this assessment to “normal worlds”—worlds where most of our general beliefs about the world (e.g. that perception is a reliable belief-forming process) are true (Goldman 1986). Now, there are a number of well-known problems for normal worlds Reliablism about epistemic

justification.²⁵ This, on its own, should set off alarm bells for externalists attracted to this way of understanding epistemic blame. But are the problems which plagued normal worlds Reliabilism also problems for normal worlds accounts of epistemic blame? There is good reason to expect they will be. If the Unwavering Externalist takes the line that subjects are epistemically blameworthy for beliefs formed by processes which tend to produce bad results in normal worlds, she'll get the result that subjects in worlds very unlike ours—for example, worlds set up such that Maud and Carl's dogmatism produces a high ratio of truths to falsehoods—are blameworthy for exercising dispositions that produce good epistemic results in their worlds. This is a surprising result about epistemic blame, and raises the question: Why should subjects in these worlds be epistemically evaluated according to dispositions that tend to produce good results in our world? This result is especially surprising if the reactive attitudes associated with epistemic blame are those considered above: reduced confidence in the subject's testimony, reduced willingness to treat them as a partner in shared inquiry, and so on. In a dogmatist-friendly world, it seems fitting to form epistemic partnerships with dogmatists. Of course, Unwavering Externalists who favor this approach may look to philosophers who have defended versions of normal worlds Reliabilism against these sorts of objections. But they should be careful to note that even if the appeal to normal worlds is defensible when it comes to understanding evaluation in terms of epistemic justification, what they need to show—in order to save the unjustified-but-blameless maneuver—is that the appeal to normal worlds is defensible when it comes to understanding evaluation in terms of epistemic blame.²⁶

²⁵ See Christensen (2007), Goldman (1986; 1988: 107-113), Graham (2017), Lemos (2007: 95-96), and Lyons (2013: 18) for discussion.

²⁶ Thanks to an anonymous reviewer for pressing me to say more here. As the reviewer points out, a defender of this view might respond as follows: The function of the epistemic reactive attitudes is to enable us to appropriately adjust our confidence in the opinions or

Alternatively, then, the Unwavering Externalist might relativize the kind of character assessment associated with blameworthiness to the subject's own world. This way of understanding blame fits better with the idea that epistemically blameless subjects are those with whom it would be fitting to maintain epistemic partnerships. But now, suppose that subjects are epistemically blameworthy for beliefs formed by processes which tend to produce bad results in their own world. This proposal also fails to secure the Unwavering Externalists desired results. On this way of understanding epistemic blame, Dupe has a terrible epistemic character: the processes which underlie her epistemic character tend to produce bad epistemic results in her world. Thus, it is not clear that accounts of epistemic blame which appeal to subjects' epistemic character can motivate the unjustified-but-blameless maneuver, especially if blameworthiness is supposed to be understood in connection to the reactive attitudes.²⁷

In this section, I considered two sets of reactive attitudes which might be relevant to evaluation in terms of epistemic blame. Neither, I concluded,

testimony of others. For this function to be fulfilled, the epistemic reactive attitudes just need to enable us to do this effectively across a broad range of circumstances in which we typically find ourselves. If we focus on the function of the epistemic reactive attitudes, it's not so implausible to think that character assessment is relativized to normal worlds. But as the reviewer also helpfully points out, once the appeal to normal circumstances is doing this much heavy lifting in the story, it's unclear why the externalist in question doesn't just go in for a straightforward normal worlds Reliabilist response to the New Evil Demon Problem. I wholeheartedly agree: It's not clear that the response in question is available to the externalists I'm targeting, who presumably want to rely on the analogy with moral blamelessness because they have independent reasons for rejecting the standard, normal worlds response to the New Evil Demon Problem.

²⁷ I think these considerations tell against at least one way of understanding Littlejohn's account of the distinction between justification and blamelessness, according to which "the agent can be excused for having ϕ -d only if she shows excellence in the exercise of her rational capacities" (Littlejohn forthcoming: 14). I will address this proposal more directly in Section 5.

can help motivate the idea that subjects like Dupe are epistemically blameless in the same way that subjects like Maud and Carl are epistemically blameworthy. Understanding epistemic blame in terms of hot reactive attitudes—like anger and resentment—gets Dupe’s case right, but Maud’s and Carl’s cases wrong. Understanding epistemic blame in terms of cold reactive attitudes—like the reduction of trust and willingness to engage in joint inquiry—gets Maud’s and Carl’s cases right, but Dupe’s case wrong. More sophisticated ways of understanding blame in terms reactive attitudes run into trouble when it comes to characterizing what it is about a subject that makes these reactive attitudes fitting. Perhaps, then, we should consider some other ways of distinguishing between epistemic justification and blamelessness.

4 Blame and Wronging

4.1 Moral Blame and Wronging

In this section, I’ll consider another way of marking the distinction between moral justification and blamelessness—one which I believe is closely related to the previous one but may be considered independently of it. The way I have in mind appeals to a distinction between actions or attitudes which are morally wrong and actions or attitudes which wrong somebody (or some group of individuals) in particular.

The distinction between wrongness and wronging is a subtle but important one. While it is not obvious that the two phenomena can come apart, there is a natural conceptual distinction to be made. Moral wrongness is the subject of countless first-order moral theories: actions or attitudes might be morally wrong in virtue of their consequences, the fact that they

violate some set of moral norms, etc. Wronging is a related, though not necessarily equivalent phenomenon. Wronging, but not wrongness per se, is supposed to involve a kind of directed relation: where there is wronging, that is, there must be one subject (or group) who wrongs, and one subject (or group) who is wronged.²⁸ By contrast, in saying that a subject acted wrongly, we needn't presuppose that this wrongfulness was directed toward any person in particular.

One way to draw the distinction between evaluation in terms of justification and evaluation in terms of blame, in the moral case, appeals to the distinction between wrongness and wronging. On this account, an action or attitude is unjustified only if that action or attitude is an instance of the kind of action or attitude which is morally wrong. A subject is blameworthy for an action or attitude, in contrast, only if she wronged somebody (or some group) via her action or attitude.

Why think that justification goes with wrongness while blame goes with wronging? The view has at least some intuitive plausibility. It is natural to think of blame as a reaction to wronging as opposed to wrongness per se. Returning to the example above, by intentionally stepping on your toe, I both do something wrong and wrong you in particular. What makes your blame appropriate, though, seems to be the wronging, not the wrongness. Another way to get a feel for this way of drawing the distinction is to consider some (admittedly controversial) cases where philosophers have thought wrongness and wronging come apart.

Here is a purported case of wrongness without wronging, due to David Owens (2012: 45-46).²⁹ Imagine I cement over the Grand Canyon,

²⁸ See Thompson (2004) and Darwall (2006) for discussion. Cases of self-inflicted wrongs, here, seem to require conceiving of one subject as two: the subject that wrongs, and the subject that is wronged.

²⁹ See Driver (2017: 6-7) and May (2015: 1) for similar cases.

reducing what was once one of America's most beautiful natural landscapes to a glorified parking lot. Intuitions here will vary, but on one reading, my action—while it is an instance of the type of action which is wrong—is not an action which wrongs anybody in particular. Again, intuitions regarding whether this really is a case of wrongness without wronging will vary. But consider how these intuitions—about whether my action merely involves wrongness or also involves wronging—pattern with intuitions about blame. If you're a person who thinks that, intuitively, this is a case of wrongness without wronging, you'll likely also have the intuition that nobody is in a position to *blame* me for cementing over the Grand Canyon (in any sense that goes beyond simply identifying me as the cause of the wrongdoing). If you're a person who disagrees—you have the clear intuition that I could and should be blamed for cementing over the Grand Canyon—it is probably because you have the intuition that I've wronged somebody (maybe the people and wildlife who could have enjoyed the Grand Canyon) in acting as I did.

In the other direction, consider a purported case of wronging without wrongness. You have just pushed a large man into the path of a trolley in order to save the five people who would have otherwise been hit by it (Thomson 1985). If this was the only way to save the five from certain death, then arguably, you wronged the large man without doing something morally wrong (Frick ms). Again, intuitions here will vary. But consider how intuitions about whether your action involves wronging without wrongness pattern with intuitions about blame. If you're convinced this is a case of wronging without wrongness, then you'll likely also have the intuition that somebody (maybe the large man's family or friends) could blame you for what you did, even acknowledging that it was the right thing to do morally speaking. If, however, you're convinced that you have not wronged the large

man, then you'll likely also have the intuition that it would be inappropriate for anyone to blame you for what you did.

Further support for these verdicts comes from thinking about the relationship between blameworthy conduct, apology, and forgiveness. Blame, as a practice, is often coupled with the practice of apology and forgiveness: blameworthy conduct warrants an apology on the part of the wrongdoer and licenses the subject who is wronged to issue her forgiveness.³⁰ And interestingly enough, intuitions about apology and forgiveness pattern with intuitions about wronging in the same way that intuitions about blame do. Starting with the Grand Canyon case: If you're a person who has the intuition that nobody was wronged by my cementing over the Grand Canyon, you'll likely also have the intuition that while I should definitely feel bad about what I did (and should take steps to try and remedy the situation), I don't owe anybody an apology for my action. On this reading, there doesn't seem to be anybody who could forgive me for my wrongdoing. In contrast, if you're convinced that I do owe somebody (or some group) an apology for my action, and think that that person (or group) is in a position to forgive me for what I did, you'll likely also have the intuition that I wronged that person (or group) by cementing over the Grand Canyon. Similar remarks apply to the trolley case. In that case, the intuition that you should apologize to the large man's family or friends, and the intuition that they can forgive you for what you did, seems to go with the intuition that you wronged him by pushing him into the path of the trolley. If you do not have the intuitions that you should apologize to the large man's loved ones or that they might forgive you for what you did, it's probably because you're convinced that he wasn't wronged by your action.

There is a conceptual connection, then, between justification and wrongness, on the one hand, and blame and wronging, on the other. This

³⁰ But see Driver (2017) for an argument against this popular view.

suggests one way of drawing the distinction between evaluation in terms of justification and evaluation in terms of blame in the moral case. Wrongness, on this picture, is a necessary condition for a subject's action or attitude being unjustified. A subject's action or attitude is justified, then, if that action or attitude was not morally wrong. Wronging is a necessary condition for a subject's action or attitude being blameworthy. A subject is blameless for her action or attitude, then, if it would not be fitting to blame her for it.

4.2 Epistemic Blame and Wronging

Can this way of drawing the distinction between justification and blamelessness help the Unwavering Externalist motivate the unjustified-but-blameless maneuver? If the distinction between epistemic justification and blamelessness is analogous to the distinction between moral justification and blamelessness with which we are already familiar, then the Unwavering Externalist should be able to characterize a notion of "epistemic wronging" relevant to our attributions of epistemic blameworthiness. But just as with the reactive attitudes approach, I think that it's unclear how Unwavering Externalists might characterize "epistemic wronging" in a way which motivates the unjustified-but-blameless maneuver.

Remember that the Unwavering Externalist wants to claim that (1) our positive intuitions about the epistemic status of Dupe's beliefs are explained by the fact that she is epistemically blameless for holding the beliefs that she does and that (2) unlike Dupe, Maud and Carl are epistemically blameworthy for holding the beliefs that they do. If epistemic blameworthiness is associated with the occurrence of "epistemic wronging", in the same way that moral blameworthiness is associated with moral wronging, then these claims have two further implications.

Starting with (1): If Dupe is epistemically blameless for holding the beliefs that she does, this should be associated with the fact that it would not be fitting to blame her for wronging anybody. Now, there are two reasons why it may not be fitting to blame a person for wronging someone. In the first case, blame is not fitting because the subject did not wrong anybody at all. In the second case, the subject did wrong someone, but should be excused or exempted from blame.³¹ So to make this implication more explicit: If Dupe is epistemically blameless for believing as she does, then it must be the case that either (a) she did not epistemically wrong anyone in holding the beliefs that she does or (b) she did epistemically wrong someone, but should be excused or exempted from blame. The implications of (2) are more straightforward: The fact that Maud and Carl are epistemically blameworthy implies that they epistemically wronged somebody in holding the beliefs that they do, without excuse or exemption.

The problem with this strategy is that there does not appear to be a notion of epistemic wrongdoing simultaneously compatible with claims (1) and (2). To see why, let's start with (1). According to the Unwavering Externalist, our favorable judgments about the epistemic status of Dupe's beliefs are judgments not about the justificatory status of these beliefs, but about the fact that they are held blamelessly. What would make these judgments apt? Again, if the notion of epistemic blamelessness is analogous to the notion of moral blamelessness with which we are already familiar, then it must be the case that either (a) she did not epistemically wrong anyone in holding the beliefs that she does or (b) she did epistemically wrong someone, but should be excused or exempted from blame. Here, option (b) seems irrelevant. Who could Dupe possibly have done wrong by, epistemically? As described, the case does not specify the existence of any

³¹ See Greco (2019) for discussion of what it might mean to have an "excuse" in the epistemic case.

other people Dupe's beliefs could have negatively impacted.³² So she must be epistemically blameless not because she has an excuse or exemption for believing as she does, but because she doesn't epistemically wrong anybody in doing so.

If this is the sense in which Dupe is epistemically blameless, however, then subjects like Maud and Carl would seem to be epistemically blameless as well. Who could Maud and Carl possibly have done wrong by, epistemically, in holding the beliefs that they do? Again, we can stipulate that Maud and Carl have not unjustly dismissed anybody's testimony, that they lack the capacity to act on their beliefs, and so on. If there is a sense in which Maud and Carl are epistemically blameworthy for holding the beliefs that they do, then, it can't be the sense in which Dupe is blameless for holding the beliefs that she does. As the cases are described, there is no one for Maud, Carl, or Dupe to epistemically wrong.

To this point, it is worth noting that most discussions of wronging in epistemology consider how subjects can wrong each other in their capacities as knowers (Fricker 2007; Marušić and White 2018) and how moral demands regarding how we relate to others can constrain what we should believe about them (Basu 2018, 2019a, 2019b). These notions of epistemic wronging, however, seem to be irrelevant to the cases we have been considering. Epistemic blameworthiness, understood in terms of a connection to epistemic wronging, does not seem to be of much use to Unwavering

³² Could the externalist say that Dupe has epistemically wronged herself by harboring a false belief? This avenue of response is certainly open to the externalist. But we should ask whether a plausible account of epistemic wronging will have the implication that one can epistemically wrong oneself merely by having a false belief, no matter how well-formed that belief is. This brings up a more general point, which I originally raised at the end of Section 2: Social epistemologists (e.g. Basu) have done a lot to develop accounts of epistemic blame and doxastic/epistemic wronging. But there is a question about whether these accounts can do the work the externalist needs to fight off the New Evil Demon problem.

Externalists who want to motivate the unjustified-but-blameless maneuver. Our comparative judgments about subjects like Dupe and subjects like Maud and Carl do not seem to track epistemic blameworthiness in this sense.

5 Blame and Responsiveness to Reasons

5.1 Moral Blame and Responsiveness to Reasons

In this section, I'll consider a third way of understanding the distinction between moral justification and blamelessness. This way of understanding the distinction is closely related to a proposal about blame which I addressed at the end of 3.2, but can be developed independently of any connection with reactive attitudes.

At the end of 3.2, I considered whether the externalist might claim that epistemic blame should be understood as a response to a subject's bad epistemic character. On this proposal, blameworthiness is associated with the fact that it would be fitting to adopt some set of reactive attitudes toward a subject on the basis of her bad character (and blamelessness with the fact that it would be unfitting to adopt such reactive attitudes). The literature on moral worth suggests a way of developing the idea that blameworthiness has something to do with a subject's character which does not rely so heavily on the connection to reactive attitudes.³³

Moral philosophers often make a distinction between an action's moral desirability and its moral worth. Arpaly writes:

³³ I thank an anonymous reviewer for pressing me on this point.

We talk about an action's moral desirability when we ask whether it is right or wrong, or how grave a wrong it is, or whether it is the best possible action (Arpaly 2003: 69).

The moral worth of an action is the extent to which the agent deserves moral praise or blame for performing the action, the extent to which the action speaks well of the agent (Arpaly 2003: 69).

Evaluating whether an action was morally desirable basically involves evaluating whether an action was morally justified, as I've been understanding justification. Evaluating an action's moral worth involves evaluating how praiseworthy or blameworthy a subject is for performing that action.

Moral desirability (and thus justification), on this picture, is characterized in terms of moral reasons: A subject's action is morally justified if and only if the action is one which her moral reasons, on balance, support. Moral worth, on this picture, is also characterized in terms of moral reasons: A subject is morally praiseworthy if and only if she performed the right action for the reasons that make it right, and morally blameworthy if and only if she performed the wrong action because she was insufficiently responsive to the relevant moral reasons.³⁴

This suggests a third way of drawing the distinction between moral justification and blamelessness. An action is morally justified if and only if that action was the right one perform, where this is understood in terms of the action's being one which the subject's moral reasons, on balance, support. A subject is morally blameless for performing an action if and only

³⁴ Arpaly (2003: 72): "For an agent to be morally praiseworthy for doing the right thing is for her to have done the right thing for the relevant moral reasons—that is, the reasons for which she acts are identical to the reasons for which the action is right."

if the subject was not insufficiently responsive to her moral reasons in so acting. Morally unjustified but blameless subjects, on this picture, are subjects who have performed an action which their moral reasons do not support, but who were not insufficiently responsive to these reasons in performing it.³⁵

5.2 Epistemic Blame and Responsiveness to Reasons

Some Unwavering Externalists favor an account of epistemic blame along these lines.³⁶ On this view: A belief is epistemically justified if and only if that belief is the right one to hold, where this is understood in terms of the belief's being one which the subject's epistemic reasons, on balance, support. A subject is epistemically blameless for believing as she does if and only if the subject is not insufficiently responsive to her epistemic reasons in so believing. Epistemically unjustified-but-blameless subjects, on this picture, are subjects who hold beliefs which their epistemic reasons do not support, but who are not insufficiently responsive to these reasons in so believing.

Can this way of drawing the distinction between justification and blamelessness help motivate the unjustified-but-blameless maneuver? The Unwavering Externalist now needs to secure the following two results:

- (1) Our positive intuitions about the epistemic status of Dupe's beliefs are explained by the fact that she is not insufficiently responsive to her epistemic reasons in believing as she does.

³⁵ In the moral case, "blameworthiness" is usually treated as a property of unjustified actions. Justified actions not performed for the right reasons are usually considered "not praiseworthy". In what follows, I will use the term "blameworthy" to cover both cases, since I think this makes the epistemic analogy a lot easier to understand.

³⁶ See Littlejohn (forthcoming: Sections 3.1-3.3) and Lord (2018).

(2) The negative correlate of these intuitions is at work in the case of BonJour's Maud and conspiracy theorist Carl. Unlike Dupe, Maud and Carl are insufficiently responsive to their epistemic reasons in believing as they do.

To repeat a now familiar refrain: I think that it's unclear how the Unwavering Externalist might characterize what it means to be responsive to one's epistemic reasons in a way which motivates the unjustified-but-blameless maneuver. To see why, consider two reasons for which Dupe might be epistemically blameless in holding some particular belief—e.g. 'There is a chair in front of me'. Either Dupe is epistemically blameless for holding this belief because (a) she correctly responds to her epistemic reasons in so believing or (b) she does not manifest a general unresponsiveness to her epistemic reasons, despite the fact that she does not in fact respond to them.

The externalist is not going to want to take option (a): that just amounts to saying that Dupe's belief is epistemically blameless because it is epistemically justified, which the externalist denies.³⁷ Dupe, then, must be epistemically blameless for believing 'There is a chair in front of me' because, while she does not respond to her epistemic reasons in holding this belief, she does not manifest a general unresponsiveness to her epistemic reasons in so believing. How should we understand this claim? The idea seems to be something like this: Dupe's belief (unfortunately) is neither true,

³⁷ To put this in terms of the moral analogy: In the moral case, responding to one's moral reasons is sufficient for justified action. But the externalist does not want to say that Dupe's beliefs are justified. Alternatively, to put this in terms of reasons: On an externalist picture of epistemic justification, responding to one's epistemic reasons involves forming a belief that meets certain externalist standards—that is true, knowledgeable, the result of a reliable belief-forming process, etc. The very idea that Dupe is unjustified-but-blameless for believing as she does presupposes that Dupe's beliefs have fallen short of such standards. So again, Dupe cannot be blameless in virtue of the fact that she's responded to her epistemic reasons.

knowledgeable, the result of a reliable belief-forming process, etc. So she does not in fact respond to her epistemic reasons in so believing. But by believing ‘There is a chair in front of me’, Dupe does not manifest a general unresponsive to the factors which would make that belief true, knowledgeable, the result of a reliable belief-forming process, etc. To put this another way: Dupe finds herself in an unfortunate situation, where the disposition to form a belief that p on the basis of it appearing as if p does not, generally speaking, lead one to form beliefs that are true, knowledgeable, reliable, etc. But put Dupe in a world where there’s no evil demon and she would respond to her reasons perfectly well. That’s the sense in which Dupe is epistemically unjustified but blameless: her dispositions are good; she’s just unlucky.³⁸

At the end of Section 3, I raised some problems for a similar explanation of Dupe’s blamelessness. I considered an Unwavering Externalists who claims that Dupe is epistemically blameless—in the sense of not being an appropriate target for reactive attitudes—because her faulty beliefs were not produced by a process which tends to produce bad epistemic results—for example, a high ratio of falsehoods to truths. The problem had to do with specifying what it means for a process to produce bad epistemic results—or more simply, what it means to have a “bad epistemic character”. Are the belief-formation processes associated with bad epistemic character ones which tend to produce bad results in all possible worlds? In “normal

³⁸ Note that Arpaly would probably not want to identify sufficient/insufficient responsiveness with such modal notions. On her view, “[t]he mere frequency or predictability of an action should not matter at all to its moral worth—unless frequency or predictability are taken to be signs of deep moral concern in the agent” (Arpaly 95). If the Unwavering Externalist wanted to characterize reasons-responsiveness without recourse to modal notions, then, she would need to develop some notion of “epistemic concern”. I think that it is very unclear what this kind of epistemic concern might involve, especially if the externalist is interested in responsiveness to epistemic reasons *de re*.

worlds”? Or in the subject’s world? None of these three ways of going, I argued, could secure the externalist’s desired result.

A completely analogous problem arises here. The Unwavering Externalist wants to say that Dupe is epistemically blameless because she is generally unresponsive to her epistemic reasons—that is, her faulty belief is not produced/maintained by a process which is generally insensitive to epistemic reasons. Now, the problem has to do with specifying what it means to be “generally unresponsive” to epistemic reasons.

If being generally unresponsive to one’s epistemic reasons means having a disposition that would, in all possible worlds, lead one to form unjustified beliefs, we get the unwanted result that Dupe (and Ruth too) is epistemically blameworthy. So the “all possible worlds” interpretation is a non-starter.

What about the “normal worlds” interpretation? If being generally unresponsive to one’s epistemic reasons means having a disposition that would, in “normal worlds”, lead subjects to form unjustified beliefs, it’s no longer clear why we should positively evaluate subjects who are reasons-responsive. To make this more concrete: Consider how Ruth and Dupe differ, in terms of their dispositions, from subjects like Maud and Carl. Ruth and Dupe, let’s suppose, both share what we might call the “anti-dogmatist’s disposition”: a disposition to revise their belief that *p* in the face of compelling evidence that *p* is false. Maud and Carl both share what we might call the “dogmatist’s disposition”: a disposition to “stick to their guns” in the face of such evidence. In normal worlds, we can suppose, the anti-dogmatist’s disposition is superior to the dogmatist’s: exercising the former, but not the latter, tends to result in beliefs that are true, knowledgeable, reliably formed, etc. Now, imagine the Unwavering Externalist claims that subjects like Maud and Carl are epistemically blameworthy because they are not sufficiently reasons-responsive: they

exercise a disposition that, in normal worlds, does not tend to result in true, knowledgeable, or reliably formed beliefs. Dupe, on the other hand, is epistemically blameless because she is sufficiently reasons-responsive: while exercising the anti-dogmatist's disposition does not tend to result in true, knowledgeable, or reliably formed beliefs in her world, the disposition would result in such beliefs in normal worlds. So far so good. But now, suppose we transported our dogmatists to a dogmatist-friendly world. There, exercising the dogmatist's disposition tends to result in beliefs that are true, knowledgeable, reliably formed, etc. If we understand reasons-responsiveness in terms of what would lead to good epistemic results in normal worlds, then our dogmatists are epistemically blameworthy in the dogmatist-friendly world. But this is a surprising result. Presumably, we care about responding to our epistemic reasons because responding to our epistemic reasons is what enables us to form beliefs that are true, knowledgeable, reliably formed, etc. If that's right, then it seems that the relevant notion of reasons-responsiveness should be relativized to the world in which we find ourselves.

This intuition might lead the Unwavering Externalist to characterize being generally unresponsive to one's epistemic reasons in terms of having a disposition that would, in the subject's own world, lead them to form unjustified beliefs. But, as we saw at the end of Section 3, this leads to the unwanted result that Dupe is epistemically blameworthy for exercising anti-dogmatist dispositions in the demon world.

To summarize: We've now examined a third way of framing the unjustified-but-blameless maneuver. This way of framing the maneuver is closely related to the reactive attitudes approach discussed in Section 3, but can be developed independently of it. As we've seen, however, understanding epistemic blamelessness in terms of reasons-responsiveness faces many of the same problems that plagued the reactive attitudes approach. So again, it's not clear that the distinction between moral justification and blamelessness

can be so easily repurposed to motivate the kind of error-theory Unwavering Externalists want to develop.

6 Rescuing the Unjustified-but-Blameless Maneuver?

Let's take stock. Unwavering Externalists need some way of explaining away the intuition that Ruth and Dupe's beliefs are equally justified. To do so, many appeal to a distinction between beliefs which are *justified* and beliefs which are *unjustified but held blamelessly*. This maneuver, if successful, would provide a simple yet powerful way of accommodating a wide range of intuitions which seem to tell against unqualified versions of externalism. Given how the familiar distinction between moral justification and blamelessness is typically understood, however, it looks like the unjustified-but-blameless maneuver cannot be carried out as neatly as proponents of the maneuver tend to assume.

There are a number of ways Unwavering Externalists might respond to the challenge as I have presented it. Before closing, I will consider two of these responses.

First, the Unwavering Externalist might respond to the challenge I've raised by objecting to the way that I've characterized the distinction between moral justification and blamelessness. This distinction, she might argue, should not be characterized in terms of reactive attitudes, the distinction between wrongness and wronging, or reasons-responsiveness; it should be understood in some other way. While this response is certainly a live option, it would, in effect, shift the burden back onto the externalist: she will need to provide a different way of distinguishing between justification and blamelessness, one which both does justice to the

theoretical roles these two forms of evaluation play in moral theorizing and secures an externalist-friendly verdict about New Evil Demon-type cases. More importantly, if I'm at least right that our *familiar* concepts of moral justification and blamelessness should be understood in one of the three ways I outlined above, then other accounts of the distinction must be at least slightly revisionary. Unwavering Externalists who object to the way that I have characterized the distinction between justification and blamelessness, then, may not be able to motivate the unjustified-but-blameless maneuver via a straightforward appeal to the moral analogy: they would need to first motivate their own substantive, potentially revisionary theory of the distinction between justification and blamelessness.

Suppose, then, that the Unwavering Externalist takes this route: she decides to take up the challenge of motivating a substantive theory of the distinction between justification and blamelessness—one which gets her the results she wants in the New Evil Demon case, albeit with significantly revisionary implications in the moral domain. This way of going, I think, faces its own challenges. To see why, consider how the Unwavering Externalist might apply Williamson's (forthcoming) framework of primary and derivative norms to the distinction between justification and blamelessness. Generally speaking, the externalist might claim, one is justified relative to a given norm if and only if one complies with that norm. One violates a norm blamelessly if and only if one complies with a derivative norm to which it gives rise (i.e. "Do as the person who tends to comply with the primary norm would!"). Here, it seems that we have a way of drawing the distinction between justification and blamelessness, independently motivated by a general—and I think, significantly revisionary—account of blame.³⁹ But why think that this is an account of *blame*? Why think that

³⁹ I say significantly revisionary because, in Williamson's framework, justification is not sufficient for blamelessness: one can comply with a primary norm while still violating the

primary norm violators, who nonetheless comply with derivative norms, are *blameless*, as opposed to having some other positive normative status? If the Unwavering Externalist wants to motivate her error-theory via the moral analogy, more needs to be said about why we should think that agents who comply with derivative norms are epistemically blameless.

To appreciate this point, consider one way the Unwavering Externalist might try to fill out Williamson’s framework. The externalist might claim that what’s good about subjects who comply with derivative norms is that they have “success-conductive dispositions”, even if they violated primary norms and so were not “successful”—a view inspired by Lasonen-Aarnio (2010, 2014). On this picture, while Dupe, Maud, and Carl all violate some primary epistemic norm, Dupe has at least believed as the person who tends to comply with that primary epistemic norm would. While this avenue of response may be available to the externalist, it would involve, in effect, abandoning the motivations behind the maneuver which I have been considering. What’s motivating the externalist’s error-theory is now the distinction between success and success-conducive dispositions, not the distinction between justification and blamelessness.⁴⁰ To repeat a point I made at the end of Section 2: I do not mean to argue that Dupe, Maud and Carl’s cases cannot be systematically distinguished via some normative distinction. What I am challenging is that this normative distinction tracks the justification/blamelessness distinction, as it is employed in the moral domain. The Unwavering Externalist might try to motivate their error-

derivative norm to which it gives rise. This strikes me as at least slightly revisionary when applied to the moral case: If one takes seriously the idea that morally justified subjects are justified for conducting themselves in a morally exemplary way, then this idea would at least seem to imply that justification entails blamelessness.

⁴⁰ As an anonymous reviewer points out, proponents of the Lasonen-Aarnio line might try to argue that these distinctions are one and the same. I agree that this is a possibility worth considering—but also that it is a substantive view which would require argument.

theory independently of the moral analogy. All I'm arguing is that I don't think this normative distinction should be understood in terms of the distinction between justification and blamelessness.

7 Conclusion

I began this paper by considering one way that Unwavering Externalists might try to explain away intuitions which appear to tell against their theories of epistemic justification. The unjustified-but-blameless maneuver, as I called it, distinguishes between two kinds of evaluation that our intuitive judgments about New Evil Demon-type cases might be tracking: evaluation in terms of epistemic justification and evaluation in terms of epistemic blame. This maneuver would seem especially well-motivated if the concept of epistemic blame relevant to the distinction between epistemic justification and blamelessness were just an extension of our familiar concept of moral blame. And this is, in fact, how Unwavering Externalists tend to motivate the maneuver. But as I have argued, it does not seem that an extension of our concept of moral blame can really help explain away the problematic intuitions that the New Evil Demon Problem elicits. This puts a heavier burden on Unwavering Externalists than is often assumed. If the concept of epistemic blame at work is not one on which we already have an independent grip, externalists cannot so easily help themselves to the unjustified-but-blameless maneuver. Such externalists, I think, will need to motivate their responses to the New Evil Demon Problem in ways which do not lean so heavily on the moral analogy.

Acknowledgements

Thank you to Evan Behrle, Dave Chalmers, David Christensen, Tez Clark, Clara Lingle, Adam Lovett, Bar Luzon, Daniel Viehoff, Crispin Wright, and especially Jane Friedman for valuable comments and discussion of earlier drafts. This paper began its life as a term paper for Ernie Sosa and Matt McGrath's graduate seminar at Rutgers, Fall 2019. I'd like to thank seminar participants, as well as audiences at NYU's Washington Square Circle and Women's Work-in-Progress Group for their constructive questions and discussion.

REFERENCES

Arpaly, Nomy. *Unprincipled Virtue: An Inquiry into Moral Agency*. Oxford: Oxford University Press, 2003.

----. *Merit, Meaning, and Human Bondage: An Essay on Free Will*. Princeton: Princeton University Press, 2006.

Basu, Rima. "Can Beliefs Wrong?" *Philosophical Topics* 46, no. 1 (2018): 1-17.

----. "The Wrongs of Racist Beliefs." *Philosophical Studies* 176, no. 9 (2019): 2497-2515.

----. "What We Epistemically Owe To Each Other." *Philosophical Studies* 176, no. 4 (2019): 915-931.

Bergmann, Michael. *Justification Without Awareness: A Defense of Epistemic Externalism*. Oxford: Oxford University Press, 2006.

Brown, Jessica. *Fallibilism: Evidence and Knowledge*. Oxford: Oxford University Press, 2018.

- "What Is Epistemic Blame?" *Noûs* 54, no. 2 (2018), 389-407.
- "Epistemically Blameworthy Belief." *Philosophical Studies* 177 (2020), 3595-3614.
- BonJour, Laurence. "Externalist Theories of Empirical Knowledge." *Midwest Studies in Philosophy* 5, no. 1 (1980): 53-74.
- Boult, Cameron. "Epistemic Normativity and the Justification-Excuse Distinction." *Synthese* 194, no. 10 (2017): 4065-4081.
- "There is a Distinctively Epistemic Kind of Blame." *Philosophy and Phenomenological Research* (2020).
- Christensen, David. "Three Questions About Leplin's Reliabilism." *Philosophical Studies* 134, no. 1 (2007):43-50.
- Cohen, Stewart. Justification and Truth. *Philosophical Studies* 46, no. 3 (1984): 279--95.
- Cohen, Stewart, and Juan Comesaña. "Williamson on Gettier Cases and Epistemic Logic." *Inquiry* 56, no. 1 (2013): 15-29.
- Darwall, Stephen. *The Second-Person Standpoint: Morality, Respect, and Accountability*. Cambridge, MA: Harvard University Press, 2006.
- Driver, Julia. "Wronging, Blame, and Forgiveness." *Oxford Studies in Agency and Responsibility Volume 4* (2017): 206.
- Frick, Johann. "Dilemmas, Luck, and the Two Faces of Morality." (manuscript).
- Fricke, Miranda. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press, 2007.
- Gardner, John. "In Defence of Defences." In *Offences and Defences: Selected Essays in the Philosophy of Criminal Law*. Oxford: Oxford University Press, 2007.
- Gerken, Mikkel. "Warrant and Action." *Synthese* 178, no. 3 (2011): 529-547.
- Goldman, Alvin. *Epistemology and Cognition*. Cambridge: Harvard University Press, 1986.
- "Strong and Weak Justification." *Philosophical Perspectives* 2 (1988): 51-69.

- Goldberg, Sanford C. "Should Have Known." *Synthese* 194, no. 8 (2017): 2863-2894.
- . *To the Best of Our Knowledge: Social Expectations and Epistemic Normativity*. Oxford: Oxford University Press, 2018.
- Graham, Peter J. "Normal Circumstances Reliabilism: Goldman on Reliability and Justified Belief." *Philosophical Topics* 45, no. 1 (2017): 33-62.
- Greco, Daniel. "Justifications and Excuses in Epistemology." *Noûs* (2019).
- Hawthorne, John, and Amia Srinivasan. "Disagreement Without Transparency: Some Bleak Thoughts 1." In *The Epistemology of Disagreement: New Essays*, by Christensen, David, and Jennifer Lackey, eds., edited by David Christensen, and Jennifer Lackey. Oxford: Oxford University Press, 2013.
- Hazlett, Allan. "Review of S. Goldberg, *To the Best of Our Knowledge: Social Expectations and Epistemic Normativity*." *Notre Dame Philosophical Reviews*, 2018.
- Hieronymi, Pamela. "The Force and Fairness of Blame." *Philosophical Perspectives* 18 (2004): 115-148.
- . "Responsibility for Believing." *Synthese* 161, no. 3 (2008): 357-373.
- Kauppinen, Antti. "Epistemic Norms and Epistemic Accountability." *Philosophers' Imprint* 18, no. 8 (2018): 1-16.
- Lasonen-Aarnio, Maria. "Unreasonable Knowledge." *Philosophical Perspectives* 24 (2010): 1-21.
- . "Higher-Order Evidence and the Limits of Defeat." *Philosophy and Phenomenological Research* 88, no. 2 (2014): 314-45.
- Lehrer, Keith, and Stewart Cohen. "Justification, Truth, and Coherence." *Synthese* 55, no. 2 (1983): 191-207.
- Lemos, Noah. *An Introduction to the Theory of Knowledge*. Cambridge: Cambridge University Press, 2007.

Levy, Neil. "The Responsibility of the Psychopath Revisited." *Philosophy, Psychiatry, & Psychology* 14, no. 2 (2007): 129-138.

Littlejohn, Clayton. *Justification and the Truth-Connection*. Cambridge: Cambridge University Press, 2012.

----- . "A Plea for Epistemic Excuses." In *The New Evil Demon*, edited by Julien Dutant and Fabian Dorsch, forthcoming.

Lord, Errol. *The Importance of Being Rational*. Oxford: Oxford University Press, 2018.

Lyons, Jack C. "Should Reliabilists Be Worried About Demon Worlds?" *Philosophy and Phenomenological Research* 86, no. 1 (2013): 1-40.

Madison, B.J.C. "On Justifications and Excuses." *Synthese* 195 (2018): 4551-4562.

Marušić, Berislav, and Stephen White. "How Can Beliefs Wrong?—A Strawsonian Epistemology." *Philosophical Topics* 46, no. 1 (2018): 97-114.

May, Simon Căbulea. "Directed Duties." *Philosophy Compass* 10, no. 8 (2015): 523-532.

McHugh, Conor. "Epistemic Responsibility and Doxastic Agency." *Philosophical Issues* 23 (2013): 132-157.

----- . "Exercising Doxastic Freedom." *Philosophy and Phenomenological Research* 88, no. 1 (2014): 1-37.

Moore, G.E. *Principia Ethica*. Cambridge: Cambridge University Press, 1903.

Owens, David. *Shaping the Normative Landscape*. Oxford: Oxford University Press, 2012.

Pryor, James. "Highlights of Recent Epistemology." *British Journal for the Philosophy of Science* 52 (2001): 95-124.

Scanlon, T.M. *What We Owe to Each Other*. Cambridge, MA: Harvard University Press, 1998.

----- . *Moral Dimensions*. Cambridge, MA: Harvard University Press, 2008.

----- "Interpreting Blame." In *Blame: Its Nature and Norms*, edited by D. Justin Coates & Neal A. Tognazzini. Oxford: Oxford University Press (2013): 84-99.

----- "Giving Desert Its Due." *Philosophical Explorations* 16, no. 2 (2013): 101-116.

----- "Forms and Conditions of Responsibility." In *The Nature of Moral Responsibility: New Essays* (2015): 89-111.

Schechter, Joshua. "No Need for Excuses Against Knowledge-First Epistemology and the Knowledge Norm of Assertion." *Knowledge First: Approaches in Epistemology and Mind*, edited by J. Adam Carter, Emma C. Gordon, Benjamin Jarvis. Oxford: Oxford University Press (2017): 132-162.

Setiya, Kieran. "Epistemic Agency: Some Doubts." *Philosophical Issues* 23 (2013): 179-198.

Sher, George. *In Praise of Blame*. New York: Oxford University Press, 2005.

Simion, Mona, Christoph Kelp, and Harmen Ghijsen. "Norms of Belief." *Philosophical Issues* 26, no. 1 (2016): 374-392.

Smart, J.J.C. "Free-Will, Praise and Blame." *Mind* 70, no. 279 (1961): 291-306.

Smith, Angela M. "Responsibility for Attitudes: Activity and Passivity in Mental Life." *Ethics* 115, no. 2 (2005): 236-271.

----- "On Being Responsible and Holding Responsible." *The Journal of Ethics*, 11 no. 4 (2007): 465-484.

Srinivasan, Amia. "Radical Externalism." *Philosophical Review* (2019).

Strawson, P. F. "Freedom and Resentment." In *Proceedings of the British Academy*, 48 (1962): 1-25. Reprinted Fischer and Ravizza 1993b: 45-66.

Sutton, Jonathan. "Stick to What You Know." *Nous* 39, no. 3 (2005): 359-396.

----- *Without Justification*. MIT Press, 2007.

Thompson, Michael. "What Is It To Wrong Someone?: A Puzzle About Justice," in *Reason and Value: Themes from the Philosophy of Joseph Raz*, edited by R. Jay Wallace, Philip Pettit, Samuel Scheffler, Michael Smith. Oxford: Oxford University Press, 2004.

Thomson, Judith Jarvis. *Rights, Restitution, and Risk: Essays, in Moral Theory*. Cambridge, MA: Harvard University Press, 1986.

Wallace, R. Jay. *Responsibility and the Moral Sentiments*. Cambridge, MA: Harvard University Press, 1994.

Watson, Gary. "Two Faces of Responsibility." *Philosophical Topics* 24, no. 2 (1996): 227-248.

Weatherson, Brian. "Deontology and Descartes's Demon." *The Journal of Philosophy* 105 no. 9 (2008): 540-569.

Williamson, Timothy. *Knowledge and its Limits*. Oxford: Oxford University Press, 2000.

-----. "Justifications, Excuses, and Sceptical Scenarios." In *The New Evil Demon*, edited by Julien Dutant and Fabian Dorsch, forthcoming.

Wolf, Susan. "Blame, Italian Style." In *Reasons and Recognition: Essays on the Philosophy of T. M. Scanlon*, by Wallace, R. Jay, Rahul Kumar, and Samuel Freeman, eds., edited by R. Jay Wallace, Rahul Kumar, and Samuel Freeman. Oxford: Oxford University Press, 2011

Zimmerman, Michael J. *An Essay on Moral Responsibility*. Totowa, NJ: Rowman and Littlefield, 1988.