

A Comprehensive Account of Blame

Self-Blame, Non-Moral Blame, and Blame for the Non-Agential

Douglas W. Portmore

Arizona State University
dwportmore@gmail.com

Blame is multifarious. It can be passionate or dispassionate. It can be expressed or kept private. We blame both the living and the dead. And we blame ourselves as well as others. What's more, we blame ourselves, not only for our moral failings, but also for our non-moral failings: for our aesthetic bad taste, gustatory self-indulgence, or poor athletic performance. And we blame ourselves both for things over which we exerted agential control (e.g., our voluntary acts) and for things over which we lacked such control (e.g., our desires, beliefs, and intentions).

Unfortunately, though, many extant accounts of blame fail to do justice to the manifest diversity in our blaming practices. For instance, T. M. Scanlon holds that “to blame a person is...to take your relationship with him or her to be modified” (2008, 128–9) and, as a consequence, “to alter or withhold intentions and expectations that that relationship would normally involve” (2013, 89). Yet, it seems clear that we can blame the dead without taking our relationship with them to have been modified in any way and without altering or withholding our intentions with respect to them. Others—e.g., Miranda Fricker (2016)—acknowledge blame's manifest diversity but hold that, given this diversity, we must give up on the prospect of providing necessary and sufficient conditions for blame. These philosophers hold that just as there is no single thing that's common to all instances of the word 'game', there is no single thing that's common to all instances of the word 'blame'. Thus, they believe that the best that we can hope for is an account that specifies the extension of 'blame' either in terms of sufficient resemblance to some paradigm or in terms of what Ludwig Wittgenstein (1953) called *family resemblances*. Still others—e.g., Angela Smith (2013)—think that although the diversity in our

blaming practices shouldn't lead us to give up on the prospect of accounting for blame in terms of necessary and sufficient conditions, we should give up on the prospect of specifying those conditions in terms of what's *constitutive* of blame. For, as these functionalists see things, the only thing that unites all instances of blame is that they all play the same functional role.¹

In contrast to all these philosophers, I'm more optimistic about the possibility of providing a set of necessary and sufficient conditions that specify blame's extension in terms of its constitution as opposed to its function. Indeed, in what follows, I'll offer a constitutive analysis of blame that can account for blame in all its disparate forms. And I'll argue that this proposal has several advantages. For one, it can account for the fact that one's having had control over whether one was to φ is a necessary condition for one's being fittingly blamed for having φ -ed. For another, it can account for why, unlike fitting shame, fitting blame is always deserved, which in turn explains why there is something morally problematic about ridding oneself of one's fitting self-blame (e.g., one's fitting guilt). Lastly, I'll show that my proposal is compatible both with the possibility that a subject lacks the standing to express her fitting blame and with the possibility that a subject ought to forgive someone whom it would be fitting for her to blame. Thus, I'll be defending what I take to be a promising, comprehensive account of blame.

1. *My Proposal for a Comprehensive Account of Blame*

To be blamed for something is to be held responsible for it. But there are at least two ways of

¹ Functionalists hold that blame is, in a certain respect, more like a mousetrap than like a diamond (Polger 2019). What makes something a mousetrap is not that it's constituted in a certain way but that it has a certain function: that of trapping a mouse. By contrast, what makes something a diamond is not that it has a certain function but that it is constituted by carbon crystals with a certain molecular lattice structure. On functionalist accounts, then, blame is just whatever has some particular role or function, and not what's constituted by certain mental states. But, like Dana Kay Nelkin (2017, 816), I doubt that our concept of blame is at bottom a functionalist one. It seems to me that some instances of blame have absolutely no function. Consider, for instance, someone who privately blames herself for some long-past misdeed only to die seconds later. I doubt that such instances of blame have a function, even an unfulfilled one. Such instances of blame don't seem defunct or deficient in the way that a mousetrap in a world of no mice does.

being responsible for something. One way is to be the cause of it. This is *causal responsibility*. Another way is to be accountable for it. And if one is accountable for something, then one can be appropriately held liable to reward or sanction for it. The reward or sanction needn't come from the law, society, or common opinion, but it must at least come from the approval or disapproval of one's own conscience—see Mill 1991, chap. 5. And, to distinguish this from causal responsibility, I'll call it *normative responsibility*.² And it's important to distinguish these two, because one can be causally responsible for something without being normatively responsible for it. For instance, I can be causally responsible for spreading a virus even if I wasn't normatively responsible for doing so given that I wasn't yet exhibiting any symptoms and, so, was completely unaware that I was infected, let alone contagious.

My aim in this paper is to provide an account of *normative blame* as opposed to *causal blame*, where to be normatively blamed for something is to be held normatively responsible for it such that one can be appropriately held to account for it. In the remainder, though, I'll typically leave the 'normative' qualifier implicit when talking about blame and responsibility. In any case, my proposed account of (normative) blame is as follows.

My Proposal: For any subject S, any potential target T (where T may or may not be identical to S), and any of T's options φ , S blames T for having φ -ed if and only if both (1) S feels disapproval of and/or disappointment in T for having φ -ed and (2) it seems to S as if T deserves to suffer the unpleasantness of guilt, regret, and/or remorse in virtue of her having violated a legitimate demand in having φ -ed.

This proposal is meant to be an account of what it is, in fact, to blame someone for having φ -ed. So, it isn't meant to be revisionary. That is, I'm not trying to figure out what blame

² The type of responsibility that contrasts with causal responsibility is more often called *moral responsibility*, but given that we can (or so I'll argue) have this sort of responsibility with respect to violations of non-moral demands, the 'moral' modifier can be quite misleading. For this reason, I've chosen to borrow Rik Peels more apt phrase *normative responsibility* (2017, 16). For someone who shares my worry about the more common phrase but adopts it anyway, see Hilary Bok (1998, 123n.1).

would need to be for our blaming practices to be justified. Indeed, I'm interested in the correct account of blame partly because I'm interested in whether our blaming practices would be justified even if all our actions are, say, causally determined. And depending on what the correct account of blame is, it will be more or less plausible to think that people can be blameworthy for acts that they were causally determined to perform. For if, on the one hand, blaming people merely entails evaluating them, then, given that evaluations can be accurate—and, thus, appropriate—regardless of whether the people being evaluated had control over the properties in virtue of which these evaluations are accurate, there would be nothing problematic about blaming people for actions that they were causally determined to perform. But if, on the other hand, blaming people entails deliberately causing them to suffer, then, given that no one deserves to suffer in virtue of things over which they lacked control, it would be problematic to blame people for acts that they were causally determined to perform—at least, it would be if we're to assume that causal determinism rules out the sort of control that's required for being deserving of suffering.

Fortunately, on my proposal, blame lies somewhere between these two extremes, such that blame goes beyond mere judgment but falls well short of necessitating the deliberate infliction of suffering.³ On my proposal, blame must go beyond mere evaluative judgment in that it necessitates feeling disapproval of and/or disappointment in its target. Thus, it requires a change in one's attitudes toward the target. And, so, there is, on my proposal, a distinction between blaming someone and merely making some set of judgments about her.⁴

But my proposal stops well short of insisting that blame must involve deliberate sanction. Since my proposal denies that blame requires taking any action, and since deliberate sanction necessitates taking action, my proposal allows that one can blame someone without deliberately sanctioning her. Indeed, on my proposal, blame essentially involves only two

³ Most agree with me in thinking that the correct account of blame must lie somewhere between these two extremes. See, for instance, Coates & Tognazzini (2013), Darwall (2010), Scanlon (2008; 2013), Sher (2006), and Smith (2013).

⁴ Thus, I concur with David Shoemaker in thinking that "blame involves attitude adjustment (and not mere deployment of judgments)" (2013, 101).

things: (1) a feeling of disapproval of and/or disappointment in its target and (2) a representation of its target as being someone who deserves to suffer the unpleasantness of guilt, regret, and/or remorse in virtue of her having violated a legitimate demand. Both are mental states as opposed to actions. Indeed, they're not even mental actions. Thus, my account allows that blame can be unexpressed and uncommunicated in that it holds that blame need only involve possessing these two mental states and, so, need not involve the act of expressing or communicating them.

And since a representation is not a belief, my proposal allows that a subject can blame someone while believing that she isn't blameworthy. So, my account allows for what might be called *recalcitrant blame* (D'Arms & Jacobson 2003), where a subject can't seem to help but blame someone despite judging that this someone isn't blameworthy.⁵ Such recalcitrant blame is analogous to recalcitrant fear, where, for instance, a subject can't seem to help but fear something (say, flying) despite judging that it poses no significant danger to her. This is possible, because just as the fear of, say, flying necessitates merely representing flying as dangerous (and, thus, not necessarily believing that it's dangerous), blaming someone necessitates merely representing this someone as deserving of the unpleasantness of guilt, regret, and/or remorse (and, thus, not necessarily believing that she deserves this). It's important to note, then, that a representation is an intellectual seeming. To represent *p* as being the case is merely for it to seem to one that *p* is the case. And just as the two lines in a Müller-Lyer illusion can seem as if they're unequal in length even if one believes that they're of equal length, it can seem to one as if *p* even if one believes that it's not the case that *p*.

Now, for a subject to represent some target *T* as deserving of *X* is for it to seem to her that, as a matter of justice and in virtue of *T*'s possessed characteristics or prior activities, *T* merits *X* in the sense that entails that the world in which *T* gets *X* and merits *X* in this sense is, other things being equal, non-instrumentally better than the world in which she gets *X* but doesn't merit *X* in this sense (cf. Feinberg 1970, 58). Note, then, that the relevant sense of 'merit'

⁵ Besides D'Arms & Jacobson, proponents of the possibility of recalcitrant blame include Carlsson (forthcoming), Gibbard (1990), McKenna (2012, 67), Pickard (2013), and Wallace (1994).

is not the one in which, say, Southwest Airlines merits a five-star customer-approval rating given its exceptional customer satisfaction. For even if Southwest Airlines does, in some sense, merit a five-star rating, it's not in the sense that entails that the world in which Southwest Airlines gets a five-star rating and merits such a rating in this sense is, other things being equal, non-instrumentally better than the world in which Southwest Airlines gets a five-star rating but doesn't merit such a rating in this sense. For if it's at all good that Southwest Airlines gets a five-star rating, it's only *instrumentally* good in that it helps customers find an airline that they'll be satisfied with. After all, there is nothing inherently good about Southwest Airlines getting a customer-approval rating that accurately reflects its degree of customer satisfaction. By contrast, there does seem to be something inherently good about someone's getting what she deserves.

But in claiming that there's something good about someone's getting what she deserves (say, X), I'm not claiming (nor am I denying) that it would be *overall* good that she gets X. Rather, I'm claiming only that the world in which she gets X and deserves X is, other things being equal, non-instrumentally better than the world in which she gets X but doesn't deserve X. Thus, to claim that someone deserves to suffer is not to claim that it would be overall good that she suffers, for the goodness of her getting what she deserves may be insufficient to compensate for the badness of the suffering that she deserves. The claim, then, is only that the world in which she suffers and deserves to so suffer is, other things being equal, non-instrumentally better than the world in which she suffers but doesn't deserve to so suffer. And it seems quite difficult to deny this claim. For, surely, we should prefer the world in which someone suffers and deserves to so suffer to the otherwise equivalent world in which she identically suffers but doesn't deserve to so suffer.

It's also important to note that blaming someone involves representing her, not as someone who deserves to suffer just any sort of unpleasantness, but only as someone who deserves to suffer the specific unpleasantness of guilt, regret, and/or remorse. What's more, blaming her involves representing her as deserving of this *in virtue of her having violated a legitimate demand* (Darwall 2013, 16). Thus, blame must involve holding its target accountable for *the violation of a legitimate demand* and not merely for *the failure to meet some normative standard*.

After all, it seems that someone can be responsible for having failed to meet a normative standard without being blameworthy.

To illustrate, consider that although my prose clearly lacks the style and elegance that you find in Oscar Wilde's writing, I can't rightly be blamed for this. For even if I could, with tremendous effort, meet such a standard, no one (not even myself) can legitimately demand that I go to such lengths given that I'm writing for those who have no right to demand such style and elegance from me: fellow academics who prefer precision to style. The thought, then, is that there are normative standards that no one (not even ourselves) can legitimately demand that we meet, and whereas we will always be blameworthy for violating a legitimate demand absent suitable excuse, we won't be blameworthy for failing to meet some normative standard unless there's someone who can legitimately demand that we meet it. Thus, the fact that I would fall short of some normative standard if I didn't ϕ may entail that I ought to ϕ , but it doesn't entail that I would be blameworthy for failing to ϕ . After all, what distinguishes that which *ought to be done because it's obligatory* from that which *ought to do be done even though it's not obligatory* is precisely the fact that only the former is conceptually tied to blame such that a failure to perform such an act without suitable excuse entails blameworthiness.⁶

Admittedly, some deny this, thinking that people can be blameworthy for performing 'suberogatory acts' (Driver 1992). Suberogatory acts are acts that are morally optional despite being *morally worse* than some morally permissible alternative. As such, they're the negative analogue of supererogatory acts—acts that are morally optional despite being *morally better* than some morally permissible alternative. The idea that suberogatory acts can be blameworthy poses a problem for my proposal given both that suberogatory acts are morally optional and that it's illegitimate to demand that an agent performs an act that's optional and, thus, not required. Of course, I think that we should just deny that suberogatory acts are ever genuinely blameworthy. But some think that it can be appropriate to resent (or otherwise blame) an agent

⁶ See, for instance, Darwall (2010, 142–43) and Portmore (2019b, 11).

for performing a suberogatory act.⁷ But, to see whether this is plausible, we should look at the sort of case that these people have in mind. Consider, then, *Kidney Donation*.

Roger and Bob are brothers. Bob is suffering from severe kidney failure. His only hope is to obtain a transplanted kidney, and the only compatible donor is Roger. If Roger donates the kidney, people respond with intense approval, because he is making a large optional sacrifice. But if he does not donate the kidney, the disapproval is also intense, even though Roger has no obligation to donate his kidney to his brother, or to anybody. Bob has no right, or entitlement, to the kidney. (Driver 1992, 287)

Now, it's clear that Roger's refusing to donate his kidney is morally worse than his agreeing to donate it. But it's not clear that his refusing to do so is morally optional and, thus, suberogatory. For it's not clear that it's illegitimate for Bob, or for others, to morally demand that he donate his kidney. Of course, I admit that Bob doesn't have the right to Roger's kidney in the sense that entitles him or others to forcibly extract it from Roger. But it can be legitimate to morally demand that someone ϕ s even if one is morally prohibited from forcing her to ϕ . What's more, it seems that whether it would be legitimate for Bob to demand that Roger donate his kidney depends, as Hallie Liberto (2012) points out, on our answers to the following sorts of questions: Is it permissible for Roger to be so partial to himself that he's permitted to choose preserving an extra kidney for himself to preserving the life of his brother? Will Bob actually die if he doesn't get Roger's kidney, or will he just be forced to live on dialysis for a few more years? What kind of relationship does Roger and his brother have? What has Bob done for Roger in the past and vice versa? It seems that we need answers to such questions before we can determine whether Roger is obligated to donate his kidney to his brother.

Of course, regardless of how we answer these questions, it will be morally best for Roger to donate his kidney. And, thus, it will be something that Roger morally ought to do. But whether donating his kidney is something that Roger is morally *obligated* to do will depend on

⁷ See, for instance, Macnamara (2013, 45). Others are less sure about whether resentment is appropriate and are confident only that anger is appropriate—see, for instance, Shoemaker (2015, 95). I concede that anger can be an appropriate response to the suberogatory, but whereas I accept that resentment is sufficient for blame, I deny that (generic as opposed to indignant) anger is.

how we answer the above questions. And it seems that if we answer them in such a way that it's intuitive to think that Roger is morally obligated to donate his kidney, then it will also be intuitive to think that he'll be blameworthy if he doesn't. And if, instead, we answer them in such a way that it's intuitive to think that Roger isn't morally obligated to donate his kidney, then it will be intuitive to think that he won't be blameworthy for refusing. So, I don't see suberogatory acts as posing a problem for my analysis of blame. As my analysis implies, it seems that people are blameworthy only for doing what's wrong and not simply for doing something morally worse than some morally permissible alternative.

So, I think that my proposal is right in insisting that blaming someone necessitates representing them as having violated a legitimate demand and not merely as having fallen short of some normative standard. But it's important to note that it's not just others that can make legitimate demands of ourselves. We can also make legitimate demands of ourselves. And, thus, we can blame ourselves for violating our own demands. Indeed, as David Shoemaker and Manuel Vargas (2019) note, "there are plenty of cases in which I may blame myself for failing to live up to ideals that I, and only I, have set for myself—for example, athletic, aesthetic, or religious ideals." And, as J. David Velleman (2003) notes, we often blame ourselves for failing to keep to some self-commitment—e.g., a commitment to maintain a certain diet or exercise regimen. So, another advantage of my proposal is that it accounts for the fact that we blame ourselves not only for failing to meet the demands of others but also for failing to meet our own demands.

And it's no accident that my proposal makes no mention of morality in specifying these demands. For, as some of the above examples illustrate, the demands that we set for ourselves needn't be moral demands. Indeed, they can be aesthetic, athletic, religious, prudential, etc. Thus, my proposal accounts for the fact that we blame ourselves not only for our moral failings, but also for our non-moral failings: for our aesthetic bad taste, gustatory self-indulgence, or poor athletic performance. And our blame of others can be non-moral as well. For consider this example from T. M. Scanlon (2013, 88): "A Mafioso can be said to blame an associate for violating the code of *omertà* [e.g., for breaking the code of silence and "ratting him out"]. Since,

we assume, the associate has good moral reason for this violation, he is not *morally* blameworthy for it... But this does not mean that his comrades' [*sic*] attitude toward him is not an attitude of blame."⁸

Lastly, my proposal allows that the target of blame can be either oneself or another, for it holds that the target of blame may or may not be identical to the blamer. And, thus, it accounts for the fact that blame can be either intrapersonal or interpersonal. What's more, it allows that the target can be alive or dead, as well as near or far away. After all, on my proposal, blaming need involve only both a feeling of disapproval and a representation of desert. And we can both disapprove of those who are dead or far away as well as represent them as having deserved (or being deserving) of the relevant sort of unpleasantness.

To sum up, I have shown, in this section, that blame requires more than mere evaluation but less than the deliberate infliction of suffering.⁹ That is, I've shown that blame involves a change in attitudes such that one must disapprove of and/or be disappointed in its target but that it needn't involve one's having the intention to make that target suffer. And I've shown that blame must involve representing its target as deserving to suffer the unpleasantness of guilt, regret, and/or remorse and that it must represent this as being so in virtue of the target's having violated a legitimate demand and not merely in virtue of her having fallen short of some normative standard. I've shown that blame needn't involve the belief that its target is blameworthy. And, thus, I've shown that blame can be recalcitrant. I've shown how, on my proposal, blame can be public or private, moral or non-moral, and intrapersonal or interpersonal. So, just in spelling out my proposal, I've already established that it can account for a lot of the diversity in our blaming practices. But there's even more that it can account for, as I'll now demonstrate.

⁸ Other proponents of the view that we can be blamed for our perceived non-moral failings include Björnsson (2017) and Matheson & Milam (2019).

⁹ In fact, it doesn't require any evaluative judgments (or other sorts of beliefs) at all. Rather, what's required is only an evaluative appraisal/representation of the sort specified in clause 2 of my proposal.

2. *What Else My Proposal Accounts for*

As George Sher and several others have noted, blame can be dispassionate such that it involves no anger, hostility, or resentment toward its target.¹⁰ “We may, for example, feel no hostility toward the loved one whom we blame for failing to tell a sensitive acquaintance a hard truth, the criminal whom we blame for a burglary we read about in the newspaper, or the historical figure whom we blame for the misdeeds he performed long ago” (Sher 2006, 88). It’s a merit of my proposal, then, that it allows for this in that one can feel disapproval without feeling any anger, hostility, or other passion. Of course, my account also allows for blame to be passionate. It is, after all, quite common to register one’s disapproval via anger and resentment. The key point, then, is that one can feel disapproval with or without anger, hostility, and resentment. And, so, my proposal allows that blame can be either passionate or dispassionate.

My proposal also allows that we can, contrary to what Scanlon claims, blame people without intending to modify our relationships with them. As Susan Wolf has noted, sometimes when we blame close family members there is a lot of screaming and remonstrance but no relationship modification (2011, 334). For, when it comes to certain family members, we have resigned ourselves to continuing on with the relationship despite everything. Of course, this doesn’t prevent us from blaming them by both disapproving of their behavior and representing them as consequently deserving of the unpleasantness of guilt, regret, and/or remorse. Indeed, our blaming them may be part of our relationship’s regular cycle in which they wrong us, we express our resentment, they express their remorse, and then we forgive them and continue on with the relationship as always.

Another merit of my proposal is that it does a nice job of accounting for the motivational tendencies associated with blame. Take self-blame. In blaming oneself for having ϕ -ed, one must suffer the unpleasantness of feeling disapproval of and/or disappointment in oneself for having ϕ -ed, which in turn involves one’s feeling guilt, regret, and/or remorse for having ϕ -ed. Such feelings are, of course, painful. And although the typical motivational tendency associated

¹⁰ See Brown (forthcoming), Smith (2013, 32), and Shoemaker & Vargas (2019).

with pain is avoidance, the main motivational tendency associated with self-blame is to seek it out both by continuing to dwell upon one's transgression and by reaching out to those whom one has transgressed so that they can further inflame one's guilt with their expressions of anger and resentment. As Herbert Morris puts it, "the man who feels guilty often seeks pain and somehow sees it as appropriate because of his guilt. ...When we think of what it is to feel guilty then, we think...of something that is owed; and pain is somehow connected with paying what one owes" (1976, 89–90).

My proposal accounts for this, because it holds that blaming oneself entails representing oneself as deserving to suffer the unpleasantness of guilt, regret, and/or remorse. Thus, in blaming oneself, it must seem that it would, to some extent, be non-instrumentally good that one suffers in this way. And this in turn explains why we're disposed to do what will inflame these feelings rather than what will extinguish them. For whereas there seems to be nothing morally problematic about our avoiding, say, headache pain, there does seem to be something morally problematic about our avoiding our appropriate guilt, regret, and/or remorse (Duggan 2018, 298). And this isn't just because headache pain isn't the sort of thing that can be appropriate. Fear can be appropriate even though there is nothing non-instrumentally good about our suffering the unpleasantness of appropriate fear.¹¹ Thus, the reason that it's non-instrumentally good to suffer appropriate guilt but not appropriate fear is that it's only the former that is deserved.¹² And it's the fact that it's deserved—and, thus, to some extent non-instrumentally good to suffer—that explains why we are often motivated to seek it out.

¹¹ Appropriate grief is a bit more complicated. There does seem to be something morally problematic about taking a pill to get rid of one's appropriate grief. But I suspect that this is because it may count as disrespectful to the one lost and/or as a form of denial that hinders one's ability to heal from that loss. In any case, the fact that there is nothing problematic about taking a pill to rid oneself of one's appropriate fear in instances where having that fear would be of no instrumental value shows that the mere appropriateness of a feeling doesn't determine whether it's something that it would be morally problematic to get rid of.

¹² Clearly, one reason that it would often be morally problematic to take a pill to alleviate one's guilt is that experiencing guilt can often be instrumentally valuable in making one less likely to commit future wrongs. Likewise, shame can be instrumentally valuable in helping one to regulate one's conduct. But it seems that it would be morally problematic to take a pill to alleviate one's appropriate guilt even if experiencing that guilt would be of no instrumental value. It would be morally problematic in that one deserves to feel bad for having violated a legitimate

Of course, another motivational tendency associated with self-blame besides this urge for a sort of self-flagellation is what Patricia Greenspan calls the “reparative urge” (1995: 130): the urge to *express* our guilt, regret, and/or remorse to those we’ve transgressed in an attempt to repair our relationship with them. And this goes hand-in-hand with one of the main motivational tendencies associated with other-blame. For one very common motivation for others in expressing their blame is to provoke or inflame guilt, regret, and/or remorse in the target of that blame. As A. P. Duggan (2018, 296) notes, expressed “blame is a form of ‘guilting’ in that blamers intend their blame to result in the blamed feeling guilty for doing wrong.”¹³ That is, blamers are moved to express their blame in the hopes that this will bring their targets both to recognize that they disapprove of what they’ve done and to share in their disapproval by coming to feel guilt, regret, and/or remorse for what they’ve done.¹⁴

So, in expressing blame, blamers deliberately aim to cause their targets to suffer. But, of course, this is potentially morally problematic, for it’s morally impermissible to deliberately inflict suffering on others unless either they deserve to so suffer or doing so is the only way to ensure a fair distribution of burdens overall. But despite being potentially morally problematic, expressions of blame actually seem to be *pro tanto* morally permissible—at least, when their targets are blameworthy (Carlsson 2017, 95).¹⁵ This means that either the blameworthy must deserve to suffer or having them suffer must be the only way to ensure a fair distribution of burdens overall. Yet it’s unclear why either would be the case? After all, to be blameworthy is just to be someone whom it is fitting to blame, where the fittingness of blame is just a matter of

demand and it is non-instrumentally good for people to get what they deserve. In this respect, then, guilt seems different from shame.

¹³ See also Carlsson (forthcoming), Fricker (2016, 167), Macnamara (2015, 559), McKenna (2012, 139–40), Wolf (2011, 338).

¹⁴ As Hannah Tierney and others have pointed out, another reason we’re often motivated to express our blame to those who have transgressed us is as a means of standing up for ourselves by expressing our sense of dignity and self-respect. See Tierney (2019), Murphy (2005, 19), and Reis-Dennis (forthcoming).

¹⁵ To say that it is *pro tanto* morally permissible for us to express our blame of the blameworthy is not to say that it is always morally permissible to do so. After all, we will, sometimes, lack the moral standing to do so. And other times it will be too risky to do so, as where the target is suicidal. I’ll have more to say about this in the next section.

its constitutive attitude—specifically, guilt, resentment, or indignation—being accurate in its representations.¹⁶ So why think that the accuracy of these representations depends upon either its target deserving to suffer or its being fair to make that target suffer? After all, it's fitting to distrust those who are untrustworthy regardless of whether they deserve to suffer the burden of being distrusted or whether it would be fair to make them suffer this burden.

Given that it's fitting to distrust the untrustworthy regardless of whether they deserve to suffer the burdens associated with being distrusted, philosophers such as Pamela Hieronymi (2004, 119–120) have argued that it is appropriate to express one's distrust of the untrustworthy even if they don't deserve to suffer the burdens associated with being distrusted, and even if there's nothing fair about their having to suffer such burdens. Now, Hieronymi admits that being the object of expressions of distrust can be just as unpleasant as being the object of expressions of blame. But she argues that this doesn't make our expressing distrust of the untrustworthy unfair or otherwise morally problematic. And, so, she concludes that there's nothing unfair or otherwise morally problematic about our expressing either distrust of the untrustworthy or blame of the blameworthy. Yet there is, I believe, an important difference between expressions of distrust and expressions of blame. Expressions of distrust don't typically aim at causing their targets to suffer. And if they did, they would be morally problematic insofar as their targets don't deserve to so suffer. And, indeed, the untrustworthy needn't deserve to suffer, for they may be untrustworthy due to no fault of their own. By contrast, expressions of blame do typically aim at causing the blameworthy to suffer, for they aim at guiltily targeting their targets. Thus, such expressions will be morally problematic unless, unlike the untrustworthy, the blameworthy necessarily deserve to so suffer. So, if we're to explain the

¹⁶ Andreas Brekke Carlsson (2017) denies that to be blameworthy is to be one whom is fittingly blamed, but one of his main motivations for doing so is that he sees no way to account for the *pro tanto* moral permissibility of expressing blame of the blameworthy on the view that the blameworthy are, most fundamentally, those who are fittingly blamed. Yet, as I now show, there is such a way. And for some of my reasons for rejecting his alternative proposal—i.e., that the blameworthy are, most fundamentally, those who deserve to feel guilty, see my (2019a). Lastly, Carlsson has so far remained neutral on what the correct general account of blame is. So, it's yet to be seen whether his alternative proposal can be combined with such an account to form a whole that's as attractive as the one that I present here.

fact that it is *pro tanto* morally permissible to express our blame of the blameworthy despite the fact that this involves intentionally inflicting suffering upon them, we must give an account of blame such that the blameworthy, unlike the untrustworthy, necessarily deserve to suffer.

It's an advantage of my account, then, that it implies that the blameworthy necessarily deserve to suffer.¹⁷ On my account, blaming a target T for having φ -ed entails representing T as deserving to suffer guilt, resentment, and/or indignation for having φ -ed. Thus, it is fitting to blame T such that T is worthy of this blame if and only if this representation is accurate. And it is accurate if and only if T deserves to so suffer. Thus, on my account, the blameworthy necessarily deserve to suffer guilt, resentment, and/or indignation. For, on my account, to be blameworthy is just to have the normative property of deserving to suffer guilt, resentment, and/or indignation. By contrast, Hieronymi holds that to be blameworthy is merely to have the descriptive property of having acted out of ill will. And, so, she thinks that a subject is blameworthy so long as that subject has in fact acted out of ill will. But given that someone can act out of ill will without deserving to suffer (consider, for instance, that young children, primitive animals, and the criminally insane can act out of ill will without deserving to suffer given their lack of free will), Hieronymi can't explain why it is *pro tanto* morally permissible to express blame even though this involves deliberately causing its target to suffer.¹⁸ So, my account has the advantage of allowing us to explain why we expect even morally good people to be motivated to express their blame with the aim of getting their targets to suffer the unpleasantness of feelings such as guilt.

Relatedly, my account of blame can explain why blame is fitting only if certain conditions are met. Take, for instance, *the control condition* (sometimes called *the freedom condition*). It holds that blaming someone for having φ -ed is fitting only if she had the relevant sort of control over whether she was to φ . My proposal can account for this so long as we assume, as seems plausible, that someone deserves to suffer some unpleasantness for having φ -

¹⁷ Many philosophers agree that the blameworthy deserve to suffer the unpleasantness of guilt, regret, and/or remorse—see, for instance, Carlsson (2017, 89) and Duggan (2018, 297).

¹⁸ I borrow this point from Carlsson (2017, 96).

ed only if she had the relevant sort of control over whether she was to φ . Thus, we can offer the following argument for the control condition.

- (P1) Blaming someone for having φ -ed entails representing her as deserving to suffer the unpleasantness of guilt, regret, and/or remorse for having φ -ed. [From clause 2 of my proposal]
- (P2) For all attitudes x , x is fitting if and only if its representations are accurate. [From the definition of 'fitting']
- (C1) Thus, blaming someone for having φ -ed is fitting if and only if she deserves to suffer the unpleasantness of guilt, regret, and/or remorse for having φ -ed. [From P1–P2]
- (P3) Someone deserves to suffer some unpleasantness for having φ -ed only if she had the relevant sort of control over whether she was to φ . [Assumption]
- (C2) Therefore, blaming someone for having φ -ed is fitting only if she had the relevant sort of control over whether she was to φ —that is, the control condition. [From C1 and P3]

And we can similarly argue for various other conditions. Take *the epistemic condition* (sometimes called *the knowledge condition*). It holds that blaming someone for having φ -ed is fitting only if she could have been reasonably expected to have known that her φ -ing would entail violating a legitimate demand. And we get an argument for this condition simply by replacing “she had the relevant sort of control over whether she was to φ ” with “she could have been reasonably expected to have known that her φ -ing would entail violating a legitimate demand” throughout the above argument.

We can also offer the following argument for *the proportionality condition*, which holds that blaming someone to degree D for having φ -ed is fitting only if D is proportional to the stringency of the demand that she violated in φ -ing (see, e.g., Fricker 2016, 168).

- (P1*) Blaming someone to degree D for having φ -ed entails representing her as deserving to suffer to degree D the unpleasantness of guilt, regret, and/or remorse for having φ -ed. [Corollary of clause 2 of my proposal]
- (P2) For all attitudes x , x is fitting if and only if its representations are accurate. [From the definition of 'fitting']
- (C1*) Thus, blaming someone to degree D for having φ -ed is fitting if and only if she deserves to suffer to degree D the unpleasantness of guilt, regret, and/or remorse for having φ -ed. [From P1*–P2]
- (P3*) Someone deserves to suffer to degree D the unpleasantness of guilt, regret, and/or remorse for having φ -ed only if D is proportional to the stringency of the demand that she violated in φ -ing. [Assumption]
- (C2*) Therefore, blaming someone to degree D for having φ -ed is fitting only if D is proportional to the stringency of the demand that she violated in φ -ing—that is, the proportionality condition. [From C1* and P3*]

We need to appeal to all three conditions in order to account for our judgments about when it is appropriate to blame people and in what degree. And accounting for such judgments is, I believe, crucial, for I agree with Scanlon that “a satisfactory account of blame should be as faithful as possible to the phenomenology of blaming and to our judgments about when it is appropriate to blame people and in what degree” (2013, 84). So, consider that without the proportionality condition we have no way of accounting for the fact that it would be inappropriate for us to blame someone who has violated a less stringent demand more harshly than we blame someone who has violated a more stringent demand. It would, for instance, be inappropriate for us to blame someone who has told a self-serving but relatively harmless lie more harshly than we blame someone who has committed murder.

We need the epistemic condition to explain why non-culpable ignorance can excuse one from being the fitting target of blame for having violated a legitimate demand. For instance, even if it's legitimate to demand that I not spread the virus that I'm carrying, it's inappropriate to blame me for doing so if I couldn't have been reasonably expected to have known that I was infected, let alone contagious.

Lastly, we need to appeal to the control condition to explain both why the only agents that we can appropriately blame are those who possess the relevant sort of control and why the only things that we can appropriately *directly* blame them for are those things over which they directly exerted such control.¹⁹ Thus, the control condition explains why newborn babies and primitive animals—both of which lack the relevant sort of control—are exempt from blame. And it explains why normal adult human beings cannot appropriately be blamed for their muscle twitches or reflex actions given that they lack the relevant sort of control over these bodily movements. What's more, it explains why a drunk driver can't be held directly responsible for failing to react in time but only for that which led to her inability to react in time—assuming that that's something over which she did exert the relevant sort of control. Perhaps, then, the only thing that we can appropriately hold her directly responsible for is her having started to drink without having first arranged for a designated driver.

Of course, some cite the fact that we often take ourselves to be (normatively) responsible for non-agential things—e.g., for desiring what's bad, believing what's contrary to the evidence, and intending to do what's incompatible with our ultimate ends—as reason for being skeptical of the control condition. But the fact that we can be responsible for such things doesn't give us any reason to doubt the control condition, but only reason to doubt that the relevant sort of control is as narrow as voluntary control. To understand why, we must understand what

¹⁹ Note, then, that I deny what's known as resultant moral luck (Zimmerman 1987): the idea that one's degree of accountability for φ -ing can be affected by the uncontrolled events that determine the results of one's φ -ing. For some compelling arguments against resultant moral luck, see Khoury (2018). And for some experimental evidence suggesting that what most affects our judgments about an agent's degree of accountability for some act is not whether, by luck, the act had a bad result but whether we judge that the agent was unjustified in believing that her act had little chance of having that bad result, see Young, Nichols, & Saxe (2010). Also, some take Frankfurt-style cases as evidence against the control condition, but see Portmore (forthcoming and 2019b) for a rebuttal.

voluntary control consists in and why we must exert it over our actions to be responsible for them.

For a subject to have *voluntary control* over an action is for her to have volitional control over whether she performs it while having rational control over whether she forms the volitions that would result in her performing it. She has *volitional control* over whether she φ s so long as, holding everything else fixed, whether she φ s just depends on whether she forms the volition to φ , and she has *rational control* over whether she ψ s (e.g., over whether she forms the volition to φ) so long as, holding everything else fixed, whether she ψ s just depends on whether and how she responds to the relevant reasons. Note, then, that volitional control over our actions is insufficient to ground responsibility for them. After all, just as I have volitional control over whether I raise my hand, a cat, presumably, has volitional control over whether it will swat at the mouse that scurries by. Yet, presumably, a cat is not responsible for swatting at the mouse because whether it forms the volition to swat isn't under its rational control. That is, whether it forms this volition is just a matter of instinct and, so, not reasons-responsive. By contrast, I can be responsible for raising my hand given that (or insofar as) whether I form the volition to do so is reasons-responsive and, thus, under my rational control. This, I've argued elsewhere, suggests that what really matters for responsibility is rational control. Indeed, it seems that the only reason that we need to have volitional control over our actions to be responsible for them is that it's only by having volitional control over our actions that we come to have rational control over them.²⁰ This is because we cannot act directly in response to our reasons; we act in response to our reasons only by being guided by our reasons to form the volitions that will, if the world cooperates, result in our acting as we so willed.

It seems, then, that we need the control condition in conjunction with the idea that the relevant sort of control is rational control to adequately distinguish between those things for which we can be held responsible—e.g., our beliefs, intentions, and voluntary actions—and those things for which we can't be held responsible—e.g., our sensations, pangs of hunger, and involuntary actions. The former are those things over which we exert rational control and the

²⁰ See also McHugh (2017, 2,749).

latter are those things over which we lack rational control. So, I admit that many of the things that we hold each other responsible for are non-agential and, thus, are things over which we lack voluntary control. But this shows, not that we should reject the control condition, but only that we should accept that the relevant sort of control is rational control. And, so, it's a merit of my proposal that it allows us to account for the fact that we can fittingly be directly blamed for the non-agential. Since, on my proposal, the variable ' φ ' ranges over all of T's *options* and not just T's voluntary *actions*, it allows that we can be fittingly blamed for our non-agential options. And, thus, we can, on my proposal, be accountable for such things as desiring what's bad, believing what's contrary to the evidence, and intending to do what's incompatible with our ultimate ends—and this is so despite the fact that we don't have voluntary control over such things.

This is important, because it seems that we can be responsible for our actions only if we can be responsible for the non-agential—specifically, for the formations of our beliefs and intentions. For as I've just shown, we can be responsible for the actions that stem from our volitions only if we're responsible for the volitions (i.e., the intentions) that gave rise to them. And, as both Nikolaj Nottlemann (2007) and Rik Peels (2017) have shown, we can be responsible for our actions and their effects only if we're responsible for our beliefs about their effects. This is because of the epistemic condition. According to the epistemic condition, one can be responsible for acting in violation of a legitimate demand only if one could have been reasonably expected to have known that so acting would constitute the violation of such a demand. To illustrate, it seems that I can be responsible for infecting my co-workers with a virus by attending a meeting only if I could have been reasonably expected to have known (and, thus, to have believed) that my doing so would infect them. So, given the epistemic condition, it seems that I can be responsible for my actions and their effects only if I'm responsible for my beliefs about their effects. And, so, if we're going to be blameworthy for anything, including our voluntary actions, the correct account of blame better allow, as mine does, for the possibility that we can be fittingly blamed for the non-agential.

Of course, many will concede that we can be fittingly blamed for the non-agential but claim that this responsibility for the non-agential must be indirect. That is, they'll appeal to the well-known tracing strategy to account for our responsibility for our forming the relevant beliefs and intentions. Now, there are, I believe, several problems with this strategy when it comes to accounting for our responsibility for our attitudes. But because many of these problems have been elucidated elsewhere, I'll mention just one in the space below.²¹

Those who employ the tracing strategy hold that someone can be responsible for, say, forming the belief that *p* even though this was never under her voluntary control because she did have voluntary control over some prior act such that she wouldn't have formed this belief had she performed (or refrained from performing) this act. So, for instance, if someone fallaciously forms the belief that taking vitamin supplements causes an increase in longevity solely on the basis of an established correlation between taking vitamin supplements and increased longevity, the tracing strategist would claim that she's responsible for forming this fallacious belief (although only indirectly) in virtue of her having been directly responsible for, say, voluntarily skipping the relevant critical thinking class—that is, the class that, had she attended, would have prevented her from making this fallacious inference. But the problem with this strategy is that it holds that what she's directly responsible for is her skipping class rather than her fallacious inference. That is, on this strategy, the demand that she is ultimately accountable for violating is, not the epistemic demand that she not infer causation solely on the basis of correlation, but the practical demand that she attend useful classes. But, intuitively, it seems that what she's ultimately accountable for is violating an epistemic demand. And this is why, when we interact with her, we're much more likely to exhort her for failing to respond appropriately to the epistemic reasons that she had for not making such an inference than we are to exhort her for failing to respond appropriately to the practical reasons that she had for attending useful classes. And this suggests that what we actually hold her accountable for is violating an epistemic demand, not a practical demand, as the tracing strategist insists.

²¹ For criticisms of the tracing strategy (where only indirect blame is appropriate for the non-agential), see Smith 2015, Vargas 2005, McKenna 2008, and Portmore forthcoming.

3. What about the Standing to Blame and the Appropriateness of Forgiveness?

I've claimed that it is *pro tanto* morally permissible to express blame of the blameworthy. But, of course, the fact that this permission is merely *pro tanto* means that it will sometimes be morally impermissible to express blame of the blameworthy. Indeed, it will be morally impermissible to do so whenever doing so would potentially have disastrous consequences, as where, say, the blameworthy target is suicidal. And it will be morally impermissible to do so whenever the blamer lacks the standing to do so.

Expressing blame of the blameworthy can be morally impermissible because of what is to be blameworthy. To be blameworthy is not to be someone to whom it is morally permissible to express blame, but rather to be someone whom it is fitting to blame, where this is a matter of the attitudes that are constitutive of blaming being accurate in their representations. Thus, expressions of blame can be inappropriate even if the associated blame is appropriate. For we must, as David Shoemaker does (2015, 221–3), draw a sharp distinction between the appropriateness of having the attitudes that are constitutive of blaming and the appropriateness of expressing those attitudes (or even acting as if one were expressing those attitudes). The latter is an overt, deliberate act, whereas the former is neither. Consequently, as Andreas Brekke Carlsson (forthcoming) points out, “the former...is matter of fittingness: whether the emotion correctly appraises its object. The latter is a matter of ethics: whether the harsh treatment is fair or deserved.”

There are many possible reasons why it could be morally impermissible for you to express your blame even if that blame is accurate in its representations. One reason is that it would be hypocritical of you to do so. Another is that it's not your place to interject yourself in this sort of private matter by doing so. Yet another is that you lack the moral authority to do so. For instance, if you've already told someone that you've forgiven her, you no longer have the authority (and, thus, the standing) to express further blame.

It can also be morally permissible to forgive the blameworthy. For even if it's fitting to blame someone, it may not be wise to do so. In this respect, blaming is like grieving. Even when

such an attitude would be correct in its representations of its object, having that attitude can be detrimental to its subject. And this explains how it can be moral and/or rational for a subject to act so to eschew blaming someone, and thereby come to forgive her, even if it would be fitting to continue to blame that someone. Thus, it's important to distinguish the fittingness of blame from the morality of acting so as to either express blame or suppress blame (and, thus, forgive). My proposal implies that whether it is fitting to blame someone is a matter of whether that someone deserves to suffer the unpleasantness of guilt, regret, and/or remorse in virtue of her having violated a legitimate demand in having φ -ed. But whether it is morally permissible to act so as to either express or suppress blame is not something that my proposal takes a stand on. So, it's a merit of my proposal that it is compatible both with the possibility that some subjects lack the standing to express even fitting blame and with the possibility that some subjects ought to forgive those whom it would be fitting to blame.

4. Conclusion

I've argued that we should accept the following account of blame.

My Proposal: For any subject S, any potential target T (where T may or may not be identical to S), and any of T's options φ , S blames T for having φ -ed if and only if both (1) S feels disapproval of and/or disappointment in T for having φ -ed and (2) it seems to S as if T deserves to suffer the unpleasantness of guilt, regret, and/or remorse in virtue of her having violated a legitimate demand in having φ -ed.

I've argued that clause 1 is necessary because blame must involve some shift in the blamer's attitudes toward the target. That shift needn't involve anger or hostility but it must at least involve disapproval and/or disappointment. And I've argued that clause 2 is necessary as well because we need it to explain our judgments about when it's appropriate to blame people and in what degree. For we need clause 2 to account for the three main conditions on blame: the control condition, the epistemic condition, and the proportionality condition. And these in turn

are needed to explain our judgments about when it's appropriate to blame people and in what degree.

What's more, I've argued that these two clauses seem to be jointly sufficient in that we don't seem to need anything beyond the two mental states specified in these two clauses to blame someone. The fact that blame can be dispassionate shows that we needn't have any passion such as anger to blame someone. The fact that we can blame people without expressing our blame and without intending to change our relationships with those whom we're blaming shows that we needn't form any intention or perform any action to blame someone. And the fact that blame can be recalcitrant shows that we needn't have any beliefs to blame someone. For instance, we can blame someone without believing that she has done wrong, manifested ill will, or is blameworthy for what she did.

Thus, I've argued that we should accept that these two clauses are individually necessary and jointly sufficient. And I've shown that this proposal allows us to account for blame in all its manifest diversity, including the fact that blame can be moral or non-moral, agential or non-agential, passionate or dispassionate, and interpersonal or intrapersonal. Lastly, I've shown that my proposal can account both for cases in which a subject lacks the standing to intentionally express her fitting blame and for cases in which she ought to forgive those whom it would be fitting to blame. So, I think that many philosophers have been too quick to give up on the possibility of providing a constitutive analysis of blame. As my proposal suggests, it seems that we can provide a set of necessary and sufficient conditions that accurately specify blame's extension and do so in terms of its constitution as opposed to its function.²²

References

²² For helpful comments and criticisms on an earlier draft, I thank Hannah Tierney.

- Bok, H. (1998). *Freedom and Responsibility*. Princeton, NJ: Princeton University Press.
- Björnsson, G. (2017). "Review of Rik Peels, *Responsible Belief: A Theory in Ethics and Epistemology*." *Notre Dame Philosophical Reviews*.
- Brown, J. (forthcoming). "What is Epistemic Blame?" *Noûs*.
- Carlsson, A. B. (forthcoming). "Shame and Attributability." In D. Shoemaker (ed.) *Oxford Studies in Agency and Responsibility*, Volume 6. Oxford: Oxford University Press.
- . (2017). "Blameworthiness as Deserved Guilt." *Journal of Ethics* 21: 89–115.
- Coates, D. J, and N. Tognazzini (2013). "The Contours of Blame." In D. J, and N. Tognazzini, (eds.) *Blame: Its Nature and Norms*, pp. 3–26. Oxford: Oxford University Press.
- Darwall, S. (2013). "Morality's Distinctiveness." In S. Darwall (ed.) *Morality, Authority, and Law: Essays in Second-Personal Ethics I*, pp. 3–19. Oxford: Oxford University Press.
- . (2010). "But It Would Be Wrong." *Social Philosophy and Policy* 27: 135–57.
- D'Arms, J. and D. Jacobson (2003). "The Significance of Recalcitrant Emotion (or, Anti-Quasijudgmentalism)." *Philosophy* 52 (Supp): 127–45.
- Driver, J. (1992). "The suberogatory." *Australasian Journal of Philosophy* 70: 286–95.
- Duggan, A. P. (2018). "Moral Responsibility as Guiltworthiness." *Ethical Theory and Moral Practice* 21: 291–309.
- Feinberg, J. (1970). *Doing and Deserving*. Princeton, NJ: Princeton University Press.
- Fricker, M. (2016). "What's the Point of Blame? A Paradigm Based Explanation." *Noûs* 50: 165–83.
- Gibbard, A. (1990). *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge, MA: Harvard University Press.
- Graham, P. (2014). "A Sketch of a Theory of Blameworthiness." *Philosophy and Phenomenological Research* 88: 388–409.
- Hieronymi, P. (2004). "The Force and Fairness of Blame." *Philosophical Perspectives* 18: 115–149.

- Khoury, A. C. (2018). "The Objects of Moral Responsibility." *Philosophical Studies* 175: 1,357–81.
- Liberto, H. R. (2012). "Denying the Suberogatory." *Philosophia* 40: 395–402.
- Macnamara, C. (2015) "Reactive Attitudes as Communicative Entities." *Philosophy and Phenomenological Research* 90: 546–69.
- . (2013). "Taking Demands Out of Blame." In D. J. and N. Tognazzini, (eds.) *Blame: Its Nature and Norms*, pp. 141–61. Oxford: Oxford University Press.
- Matheson, B. and P.-E. Milam (2019). "The Case for Non-Moral Blame." Manuscript.
- McHugh, C. (2017). "Attitudinal Control." *Synthese* 194: 2,745–62.
- McKenna, M. (2012). *Conversation and Responsibility*. Oxford: Oxford University Press.
- . (2008). "Putting the Lie on the Control Condition for Moral Responsibility." *Philosophical Studies* 139: 29–37.
- Mill, J. S. (1991). [1861]. *Utilitarianism*. Reprinted in J. M. Robson (ed.), *Collected Works of John Stuart Mill*. Vol. 10, pp. 203–59. London: Routledge.
- Morris, H. (1976). *On Guilt and Innocence: Essays in Legal Philosophy and Moral Psychology*. Berkeley, CA: University of California Press.
- Murphy, J. (2005). *Getting Even: Forgiveness and Its Limits*. New York: Oxford University Press.
- Nelkin, D. K. (2017). "Blame." In K. Timpe, M. Griffith, and N. Levy (eds.) *The Routledge Companion to Free Will*, pp. 374–88. New York: Routledge.
- Nottlemann, N. (2007). *Blameworthy Belief. A Study in Epistemic Deontology*. Dordrecht, Netherlands. Springer.
- Peels, R. (2016). *Responsible Belief: A Theory in Ethics and Epistemology*. Oxford: Oxford University Press.
- Pickard, H. (2013). "Irrational blame." *Analysis* 73: 613–26.
- Polger, T. W. (2019). "Functionalism." *The Internet Encyclopedia of Philosophy*, ISSN 2161-0002, <https://www.iep.utm.edu/>. Accessed on May 21, 2019.

- Portmore, D. W. (forthcoming). "Control, Attitudes, and Accountability." In D. Shoemaker (ed.) *Oxford Studies in Agency and Responsibility*, Volume 6. Oxford: Oxford University Press.
- . (2019a). "Desert, Control, and Moral Responsibility." Manuscript.
- . (2019b). *Opting for the Best: Oughts and Options*. New York: Oxford University Press.
- Reis-Dennis, S. (forthcoming). "Anger: Scary Good." *Australasian Journal of Philosophy*.
- Scanlon, T. M. (2013). "Interpreting Blame." In D. J. and N. Tognazzini, (eds.) *Blame: Its Nature and Norms*, pp. 84–99. Oxford: Oxford University Press.
- . (2008). *Moral Dimensions. Meaning, Permissibility and Blame*. Cambridge: Harvard University Press.
- Sher, G. (2006). *In Praise of Blame*. Oxford: Oxford University Press.
- Shoemaker, D. (2015). *Responsibility from the Margins*. Oxford: Oxford University Press.
- . (2013). "Blame and Punishment." In D. J. and N. Tognazzini, (eds.) *Blame: Its Nature and Norms*, pp. 100–18. Oxford: Oxford University Press.
- Shoemaker, D. and M. Vargas (2019). "Moral Torch Fishing: A Signaling Theory of Blame." Manuscript.
- Smith, A. M. (2015). "Attitudes, Tracing, and Control." *Journal of Applied Philosophy* 32: 115–32.
- . (2013). "Moral Blame and Moral Protest." In D. J. and N. Tognazzini, (eds.) *Blame: Its Nature and Norms*, pp. 27–48. Oxford: Oxford University Press.
- Tierney, H. (2019). "Guilty Confessions." Manuscript.
- Vargas, M. (2005). "The Trouble with Tracing." *Midwest Studies in Philosophy* 29: 269–90.
- Velleman, J. D. (2003). "Don't Worry, Feel Guilty Royal Institute of Philosophy Supplement 52: 235–48.
- Wallace, R. J. (1994). *Responsibility and the Moral Sentiments*. Cambridge: Harvard University Press.

- Wittgenstein, L. (1953). *Philosophical Investigations*. G. E. M. Anscombe and R. Rhees (eds.), G. E. M. Anscombe (trans.). Oxford: Blackwell.
- Wolf, S. (2011). "Blame Italian Style." In R. J. Wallace, R. Kumar, and S. R. Freeman (eds.), *Reasons and Recognition: Essays on the Philosophy of T. M. Scanlon*, pp. 332–47. New York: Oxford University Press.
- Young, L., S. Nichols, and R. Saxe. (2010). "Investigating the Neural and Cognitive Basis of Moral Luck: It's Not What You Do but What You Know." *Review of Philosophy and Psychology* 1: 333–49.
- Zimmerman, M. J. (1987). "Luck and Moral Responsibility." *Ethics* 97: 374–86.