

SHAME AND THE SCOPE OF MORAL ACCOUNTABILITY

BY SHAWN TINGHAO WANG

Abstract

It is widely agreed that reactive attitudes play a central role in our practices concerned with holding people responsible. However, it remains controversial which emotional attitudes count as reactive attitudes such that they are eligible for this central role. Specifically, though theorists near universally agree that guilt is a reactive attitude, they are much more hesitant on whether to also include shame. This paper presents novel arguments for the view that shame is a reactive attitude. The arguments also support the view that shame is a reactive attitude in the sense that concerns moral accountability. The discussion thereby challenges both the view that shame is not a reactive attitude at all, suggested by philosophers such as R. Jay Wallace and Stephen Darwall, and the view that shame is a reactive attitude but does not concern moral accountability, recently defended by Andreas Carlsson and Douglas Portmore.

Keywords: shame, guilt, reactive attitude, moral accountability, moral attributability, moral responsibility

1. Introduction

There are a variety of negative emotions a person may experience in the face of a moral infraction: resentment, indignation, guilt, fury, fear, disgust, shame, and so on. Philosophers often categorize these emotions into two groups: some, including at least resentment, indignation, and guilt, are *reactive attitudes*; others, such as fury and fear, are *non-reactive attitudes*. The labels are attached in terms of each emotion's relation to moral responsibility. Reactive attitudes, unlike non-reactive attitudes, target participants of the practices involving holding each other, and holding oneself, morally responsible. The propriety conditions of reactive attitudes, unlike those of non-reactive attitudes, require that their target in fact bears moral responsibility. Though non-reactive attitudes may serve important functions, according to the standard philosophical view, it is the reactive attitudes that play the most central role in those activities involving holding people morally responsible.

Philosophers debate the scope of reactive attitudes, and it is particularly controversial where *shame* fits in. Is shame a reactive attitude? Though Strawson's (1962) seminal paper on reactive attitudes endorses a positive answer, the view has been frequently challenged. Wallace (1994) restricts reactive attitudes to only resentment, indignation, and guilt, but excludes shame, because shame is not constitutively linked to 'holding oneself to a demand' (241) as guilt is. Darwall (2006) also excludes shame in his second-personal framework, arguing that shame inhibits the kind of 'second-personal engagement' (72) that is, according to Darwall, central to the interpersonal functional role of reactive attitudes. The assumption that excludes shame as a reactive attitude also seems implicit in many other

theorists' work, given the disproportionate attention they choose to devote to guilt over shame. Just as one example, Clarke (2013, 2016) offers a defense of retributivism based on the idea that responsible wrongdoers deserve to feel guilt, but he very quickly dismisses the relevance of shame, claiming that it is not 'a matter of justice' (2016: 128) that one feels shame about failing to live up to certain standards.

Another possible answer goes beyond the simple 'yes' or 'no', but rather appeals to a distinction between different senses of reactive attitudes according to their connections to different senses of moral responsibility. Watson (1996) has distinguished between *attributability* and *accountability* concepts of responsibility. Here is a rather minimal way to understand the two concepts, when applied to things that are morally bad or morally wrong, namely, things that we are *negatively* morally responsible for: A person is morally responsible for *X* in the attributability sense just in case *X* reveals her evaluative commitments, namely, what she values. Sometimes attributability is defined in terms of whether *X* reflects one's substandard moral character or traits. I am inclined to think the definition in terms of evaluative commitments is more plausible, but I expect the arguments and views in the current paper to apply to the character-based definition of attributability just as well as it applies to the value-based definition. By contrast, a person is negatively morally responsible for *X* in the accountability sense just in case *X* provides a pro tanto moral reason (or at least a pro tanto reason that is understood within our practices to be moral) for other people to

hold resentment or indignation toward her.¹ It is also commonly held that, if a person is negatively morally responsible for *X* in the accountability sense, then she is *morally blameworthy* for *X*, which further entails that there is a pro tanto moral reason (or at least a pro tanto reason that is understood within our practices to be moral) for other people to *blame* her.² This does *not* always give rise to an all-things-considered reason to blame, because there may be contexts where the pro tanto moral reason is overridden, e.g. when the potential blamer lacks the standing to blame. Sometimes accountability is defined in terms of more substantial conditions such as whether the person deserves sanction in virtue of *X* or whether *X* violates an interpersonal demand. I will proceed with the more minimal formulation of accountability, and address these more substantial views about accountability in section 4. On the basis of the distinction between attributability and accountability, Carlsson (2019) and Portmore (2019a, 2019b) defend an intriguing position that shame is a reactive attitude but it is concerned with the attributability rather than accountability sense of moral responsibility; guilt, they suggest, is concerned with accountability instead. Hence, we have three available answers to the question raised in the beginning: *exclusivism*, the view that shame is not a reactive attitude at all; *inclusivism*, the view that shame is a reactive

¹ This involves some simplification. Some have thought that accountability only entails a *conditional* pro tanto reason to hold the negative reactive attitudes (see Nelkin 2016).

² Theorists disagree on what blame is. On my view, blame at least includes resentment and indignation (blaming attitudes) and the outward manifestations of resentment and indignation (blaming behaviors). But I will stay neutral about whether there can be other forms of blame as well. Although I will not try to argue for it here, I am inclined to think that the mere judgment of a fault in one's conduct—that is, what Watson calls 'aretaic blame' (1996: 238)—is not yet sufficient to count as blame.

attitude in the sense that concerns moral accountability; and *divisionism*, the view that shame is a reactive attitude but not in the sense that concerns moral accountability.

In this paper, I aim to advance novel arguments for inclusivism, what has now become the heterodox view in the literature.³ Throughout the paper, I will assume a commonly accepted picture of reactive attitudes: their *targets* are the good will, ill will, or indifference of moral agents as participants of those practices involving holding each other, and holding oneself, morally responsible; and their *propriety conditions* are as follows: it is appropriate to hold a reactive attitude toward a person just in case the person is morally responsible for something. We can then distinguish between *attributability-tied reactive attitudes* and *accountability-tied reactive attitudes*. The former kind of attitudes is appropriate toward a person just in case the person is morally attributable for something, whereas the latter kind of attitudes is appropriate toward a person just in case the person is morally accountable for something. There has been a debate about what exactly the appropriateness consists in, and whether it means the same relation for attributability-tied and accountability-tied attitudes (e.g. Shoemaker 2017; Carlsson 2019; Portmore 2019a, 2019b). I will not take a stand on this controversy. Rather, my approach is to start from two assumptions that I believe different parties in the debate would agree on. First, the propriety condition for attributability-tied reactive attitudes is the *fitting condition* of the

³ There are important exceptions. Watson (2014) offers a brief but insightful defense of inclusivism. Shoemaker distinguishes between different senses of responsibility, but claims what he calls ‘agential shame’ in fact ‘cuts across all categories’ of moral responsibility, including attributability, answerability, and accountability (2014: 26).

emotion, which further involves having correct representations. Note that this alone does not ensure that we have any moral reasons to hold the emotion, since many ordinary fitting conditions do not entail any moral reasons at all, for example, when a joke is fitting because it is funny. Second, the propriety condition for accountability-tied reactive attitudes should provide *moral reasons*, or at least reasons that are understood within our practices to be *moral*, to hold the attitude. This is compatible with both the claim that it is not a fitting condition at all (Carlsson 2019) and the claim that it is still a fitting condition but gives rise to moral reasons due to the fact that the content of the accountability-tied attitudes is always about desert, a moral-reason-giving notion (Portmore 2019a, 2019b).

Another preliminary point concerns what it exactly means to say that shame is a reactive attitude. Does it mean some particular *tokens* of shame are reactive, or the general emotional *type* is as such? I am inclined to endorse both. There are some particular tokens of shame as accountability-tied reactive attitudes, and, as a further point, those tokens belong to an emotion type that we can properly refer to as an accountability-tied reactive attitude. It is worth noting that, however, we may have to do some conceptual regimentation to single out this type of shame, perhaps as ‘moral shame’ or ‘agential shame’, in order to exclude, for instance, physical shame about one’s bodily features—which intuitively falls outside the domain of reactive attitudes. But another possibility is that shame is a kind of accountability-tied reactive attitude, even though it has some subcategory, such as physical shame, that we rarely have moral reason to hold and is thereby rarely appropriate. This is a choice point, I think, that ultimately depends on how we individuate psychological kinds, a difficult task that I will leave aside. Rather, my thesis is

that either shame or a subcategory of shame is an accountability-tied reactive attitude. If the thesis holds, this would be a success for the inclusivist, since the current defenders of exclusivism or divisionism apparently think neither shame nor a subcategory of shame counts as an accountability-tied reactive attitude.

2. Shame and Guilt: Intertwinedness and Differentiation

Here is a first attempt to develop an argument for inclusivism. Note that both exclusivism and divisionism presuppose that shame and guilt are quite distinct. But this appears to be in tension with the phenomenology that shame and guilt are deeply intertwined with each other. On my usage of the term, two emotions are deeply intertwined with each other when, in a wide enough range of cases, they are *overlapped*, *undifferentiated*, and *convertible to each other*. First, overlapping means that, in many cases, the two emotions are both experienced by a person and often come and go together. Second, they are undifferentiated in the sense that, in many cases, a person does not differentiate between the two emotions, and, upon reflection, it can be hard for a person to tell which of the two emotions she feels. Finally, they are convertible to each other such that intense feelings of one emotion sometimes shift over and become the other, and/or the other way around. Shame and guilt satisfy all three conditions, as supported by both reflection on ordinary experiences and empirical evidence.

(i) Though shame and guilt can come apart, they significantly overlap in a large number of cases. This is further supported by experiments showing that only very few antecedents uniquely elicit either shame or guilt, whereas a large majority of antecedents elicit both

(Keltner & Buswell 1996; Tangney 1992). (ii) We can find it quite hard to distinguish if our emotion is shame, guilt, or both, at least for many cases. From college students to clinical psychologists, as Tangney and Dearing (2002) observe, people use ‘shame’ and ‘guilt’ interchangeably and perhaps even inconsistently. It is not uncommon at all for people to talk in terms of the phrase ‘shame and guilt’ and discuss their sources and effects together, with little discrimination between the two emotions. (iii) Guilt experiences appear to be able to be converted to shame experiences (e.g. Lewis 1971). These features involving intertwinedness may appear to pose a prima facie challenge against exclusivism and divisionism, thereby supporting inclusivism. That is, since (1) guilt is an accountability-tied reactive attitude, and (2) shame and guilt are deeply intertwined, we should conclude that (3) shame is also an accountability-tied reactive attitude. I will refer to this as the *Intertwinedness Argument* for inclusivism.

Though I will ultimately contend that the intertwinedness argument, when suitably revised and supplemented, can pose a strong case for inclusivism, the inference as stated in the above form is unconvincing and subject to counterexamples. One may think resentment is deeply intertwined with what Pereboom calls ‘fury’, a more primitive kind of anger that we share with ‘bears and wolves’ (2014: 147). But resentment is, whereas fury is not, a reactive attitude. More generally, the inference seems to overlook the distinction that Wallace (1994) makes between *natural* attitudes and *reactive* attitudes. Wallace argues that inclusivists make the mistake of confusing the two and thinking whatever attitudes that are natural to have in interpersonal relationships must be reactive. He then suggests that emotions like shame are natural but non-reactive attitudes. Similarly, Wallace might reply

that intertwinedness shows that both emotions are natural, but not that they should be similarly treated when it comes to categorizing emotions as attributability-tied reactive, accountability-tied reactive, or non-reactive.

Nonetheless, I take it that the intertwined argument can still be modified into a forceful, burden-shifting move in support of inclusivism. This is because deep intertwinedness between the two emotions should at least shift the burden to those who reject inclusivism to show how differences between shame and guilt justify treating only the latter as an accountability-tied reactive attitude. This burden is met in the comparison case of the difference between fury and resentment. In contrast to resentment, fury involves a distinct self-defense mechanism that registers the information that ‘there’s a threat to be violently neutralized’ (Pereboom 2014: 147), which plays little role in our responsibility practices. The task falling upon us then is to investigate if there are similar reasons to differentiate shame from guilt in a manner that justifies categorizing them differently.

Just as we can differentiate fury from resentment despite their intertwinedness, there are ways to differentiate shame from guilt too, in terms of a cluster of features including their phenomenology, representational content, and motivational tendency. Here is a rough characterization of the major differences in these three regards. The phenomenology of shame is an intense painful feeling that can be said to involve a wish to hide or disappear, fused with a disruption in thought, a sense of confusion, and sometimes blushing (Lewis 2008: 748); the representational content of shame is a failure to live up to some norm or expectation, where the failure is primarily concerned with one’s substandard self (Niedenthal, Tangney, & Gavanski 1994; Tracy & Robins 2006; Tangney, Stuewig, &

Mashek 2007; Lewis 2008); the motivational tendency is to dissipate the emotion by ‘reinterpretation, self-splitting, or forgetting’ (Lewis 2008: 748). In contrast, the phenomenology of guilt is a painful feeling, but generally less intense than shame (Tangney 1995); the representational content of guilt is a particular transgression typically in the form of a norm-violating action (Niedenthal et al. 1994; Tangney et al. 2007; Lewis 2008); the motivational tendency is to direct the agent to perform a corrective action to repair the failure (Lewis 2008: 748).

Many details of this characterization are controversial and up for debate, but I take it to be a good starting point to consider those familiar ways in which shame and guilt can come apart. The characterization also fits well with empirical findings. In particular, the self-action contrast is widely accepted in social psychology: shame is more concerned with failure that reflects one’s substandard *self*, whereas guilt is more concerned with failure in a *particular action*. Tangney, Stuewig, & Mashek put it that shame is ‘a negative evaluation of the global self’, whereas the focus of guilt is ‘a negative evaluation of a specific behavior’ (2007: 349). Shame brings the agent to see a shortcoming in herself, to think ‘if only I weren’t ...’, while guilt brings the agent to see a failure in her specific behavior, to think ‘if only I hadn’t ...’ (Niedenthal et al. 1994). This is ‘the currently most dominant basis for distinguishing between shame and guilt’ (Tangney et al. 2007: 349) in the psychological literature. In philosophy, many draw a broadly similar self-action distinction between shame and guilt as well. For example, Rawls argues shame and guilt have different explanations, using the example of a person who ‘cheats or gives in to cowardice’ and feels both guilt and

shame (1971: 391). According to Rawls, this person's guilt is due to the quality of his action, but his shame is due to a diminished sense of 'his own worth' (1971: 391).

Let me stress that the self-action contrast does not require the target of shame to be only about the self or the target of guilt to be only about actions. We can feel both shame and guilt toward an action, but, the self-action contrast suggests, shame about the action is more focused on the ways in which the action reflects on the person's self, whereas guilt about the action focuses on the qualities of the transgressing action itself as the locus of the emotion. Shame can sometimes be about very local failures as well, for example, when a person is ashamed of her fleeting bad thought. However, such shameful thoughts usually accompany a further judgment like the following: 'How can *I* have such a thought?' If no such judgment would arise for the person, then, I think, we have good reason to doubt if her emotion is shame at all. That is, the target of shame seems to be typically (though perhaps not always) wrapped up in some of our general character traits or dispositions, even when the same object can also be the target of guilt.

Given that the self-action contrast is the most dominant basis for differentiating shame from guilt, the most natural rejoinder to the intertwinedness argument is to use this contrast against inclusivism. This is surely not the only possible rejoinder. Many philosophers have appealed to other ways of contrasting shame with guilt in order to classify only guilt as an accountability-tied attitude. For example, Wallace (1994) argues that guilt rather than shame is involved in holding oneself to an expectation or demand; and Darwall (2006) suggests that guilt rather than shame exhibits a 'second-personal' standpoint. These views usually come with a fairly substantial position about why moral accountability is a

significant concept for us. I will come back to some of these approaches later and address the significance of moral accountability in section 4.

Now consider the following line of argument. (1) The primary object of concern in guilt, but not in shame, is the particular action that constitutes wrongdoing. (2) The only things we are ever directly morally accountable for are particular actions that constitute wrongdoings. But (3) an accountability-tied reactive attitude must have the same primary object of concern as what we are directly morally accountable for. Therefore, (4) shame cannot be an accountability-tied reactive attitude, but guilt can. Call this the *Action-Based Objection* to inclusivism. The thought is that, if moral accountability has particular actions as the primary focus, then we should think only the emotion that has particular actions as its primary focus—guilt but not shame—is an accountability-tied reactive attitude. I think the objection has some initial plausibility. In particular, the assumption that we are only ever directly accountable for actions is an intuitive and traditional view. After all, how can a person be morally accountable if not due to something that she *did*? How can a person be morally blameworthy if she has not *done anything wrong*? It appears to be a plausible view that we are only *indirectly accountable* for things other than actions, when those things are the reasonably foreseeable consequences of what we do.

But why restrict the scope of moral accountability in this way? One common answer is that only actions are *within our voluntary control*, in the sense that we can *choose* whether to perform some actions. This connects to another, related argument against inclusivism, directly from the assumption that voluntary control is a necessary condition for moral accountability. The argument goes as follows. (1) Guilt is appropriate only if its target is

within the person's voluntary control, but shame can be appropriate even if its target is beyond the person's voluntary control. However, (2) the only things we are ever directly morally accountable for are those within our voluntary control. Given that (3) an accountability-tied reactive attitude must have its appropriate condition being restricted by the necessary conditions of direct moral accountability, it follows that (4) shame cannot be an accountability-tied reactive attitude, but guilt can. Call this the *Control-Based Objection* to inclusivism. Carlsson (2019) and Portmore (2019a, 2019b) have both discussed a similar line of argument to defend divisionism.⁴ Again, I find the objection has strong initial appeal. (1) is backed up by both intuitive cases and empirical results. We frequently feel shame about our beliefs, desires, or characters, which do not seem to be what we can directly choose to alter. Nonetheless, it appears that shame about these targets can be appropriate. Tracy and Robins (2006) presented evidence that 'uncontrollable attributions for negative events' lead to greater shame than guilt, while 'controllable attributions for negative events' lead to greater guilt than shame (1347). (2) is a traditional and widely accepted picture about accountability: we are morally accountable for things that we voluntarily choose. It has been suggested that only such things can make blame *fair* and can make us *deserve* being blamed and feeling painful emotions. (3) also appears to be a reasonable constraint on the accountability-tied reactive attitudes.

⁴ Carlsson appears to endorse the argument, but he does not specify if the relevant notion of control is voluntary control. Portmore explicitly claims that the argument only works if control means something weaker than voluntary control.

Nonetheless, I think both the *Action-Based Objection* and *Control-Based Objection* should ultimately be rejected. Though they make sense on the assumption that direct accountability is only about actions or that direct accountability is only about things within our voluntary control, both assumptions should be rejected. Recent work on moral responsibility has offered reasons for expanding the scope of direct moral accountability beyond particular actions and beyond the domain of voluntary control. A closer examination on these developments, I shall argue, in fact lends support to treating shame as an accountability-tied reactive attitude. Just as an important facet of our practices of holding people morally accountable is more action-directed and thereby more strongly associated with guilt, these practices have another important facet that is more self-directed and thereby more strongly associated with shame.

3. Expanding the Scope of Moral Accountability

There have been prominent defenses of the claim that direct moral accountability does not just involve actions or just involve things within our voluntary control. Theorists argue that the scope of direct moral accountability should be expanded to include, for example, *attitudes*, *character*, and *the self*.⁵ I contend that, if these defenses are on the right track, then

⁵ See, e.g. Smith (2005, 2012), Sher (2006), Hieronymi (2008), Graham (2014), Adams (1985), Holroyd (2012), Taylor (1976), Westlund (2003). Though some of these authors do not use the term 'moral accountability', I think this is a fair representation of their view. For example, Sher, Hieronymi, and Graham talk in terms of blameworthiness, which is typically associated with accountability; though Smith (2005) talks primarily in terms of attributability, she extends the idea to accountability in her more recent work (2012).

the *Action-Based Objection* and the *Control-Based Objection* are undermined, and there are instead good reasons to endorse inclusivism.

To start with, consider three cases as follows:

Sexist Attitude. Amanda is a physics professor. When she teaches Physics 101, she consciously holds the belief that, in her class, male students are much more talented than female students in physics. But she neither expresses this sexist belief nor performs any discriminatory action based on this sexist belief.

Cold Character. Barbara has a character of cold-heartedness. For example, she has strong dispositions to dismiss her friends' emotions as unimportant, to not care about her close family members, and to not sympathize with those people who suffer from hard circumstances. But she neither expresses this cold character nor performs any cold action based on this character.

Cruel Self. Claire is a cruel person, and, furthermore, she identifies with her cruelty. She desires to see innocent people suffer, and she wants this desire to be effective in her actions. But she neither expresses her cruel self nor performs any cruel action based on her cruel self.⁶

⁶ Here I assume a view like Frankfurt's (1971), according to which one's self involves one's second order volitions, that is, those desires one desires to have and to be motivated by. This is not essential to my

In all three cases, the agent has a substandard attitude, character, or self without further performing any substandard action. Many details of the cases are left open, so the agents surely can have some excuses that help them escape being blameworthy. But can there be *some* contexts where they are morally accountable for the attitude, character, or self, in the absence of morally accountable actions? That is, does the fact that they did not perform any actions on the basis of the attitude, character, or self *exclude* the possibility of accountability? My guess is that people's intuitions will be varied. Laypeople's pretheoretical intuitions will most likely accompany a sense of confusion: it may initially seem that Amanda, Barbara, and Claire can surely be blameworthy, but the fact that they have not performed any substandard actions may also generate hesitation for people's having any clear-cut pretheoretical intuitions about the cases.

Here is a possible reaction upon reflection: Amanda, Barbara, and Claire are blameworthy, but only *indirectly* so. That is, they are blameworthy for the attitude, character, or self only to the extent that they were blameworthy for certain actions they did before that have the attitude, character, or self as reasonably foreseeable downstream consequences. This is a reasonable implication to draw from the cases, and I will consider this line of thought later in section 5.1.

My focus, however, will be on another kind of diagnosis, according to which Amanda, Barbara, and Claire can be *directly morally accountable*, and thereby directly

argument, and I expect that similar cases and arguments can be developed even if one adopts alternative reasonable views of the self.

blameworthy for the attitude, character, or self. Theorists have provided various reasons to support this kind of diagnosis. The first is an argument based on ordinary practices.

Regardless of our intuitions about the above cases, as theorists observe, we do in *our practices* blame people for their attitudes even in the absence of blameworthy actions, including ‘attitudes like fear, contempt, admiration, guilt, envy, and resentment’ (Smith 2005: 254), beliefs and judgments (Hieronymi 2008), self-centered attitudes (Adams 1985), disrespect (Graham 2014), and even implicit attitudes (Holroyd 2012); character and traits (Sher 2006); and a person’s self (Taylor 1976; Westlund 2003). These blaming practices, just as blaming practices based on actions, can be in either emotional forms, like resentment and indignation, or behavioral forms, like the outward manifestations of resentment and indignation. Though our blaming practices are fallible evidence for blameworthiness, they are evidence nonetheless, and a picture that renders inappropriate this wide range of blaming practices is hard to defend. The second is an argument based on quality of will, which plays a particularly important role in the defenses of responsibility for attitudes by Smith (2005) and Hieronymi (2008). They contend that attitudes can indicate a flaw in the extent to which one has *proper regard* for another person’s moral standings and interests. This, Smith and Hieronymi suggest, is a sufficient basis for holding resentment or indignation toward people like Amanda. A similar line of argument can be extended to cases involving character and self too: to the extent that they indicate, or even constitute, a substandard quality of will, they are the proper targets of moral accountability. A third, related argument concerns the connection between responsibility for action and responsibility for attitude, character, and self. Sher (2006) highlights this tight connection

by pointing to scenarios where an action's badness is a 'near-inevitable' result of her bad character—which, according to Sher, suggests that she should be blamed for the character as well (2006: 65). The argument can be extended to cases involving attitudes and self: there are scenarios where an action's badness is a near-inevitable result of one's substandard attitudes, or one's substandard self, and, as a result, it will appear arbitrary to *categorically* exclude attitudes and self from the domain of direct moral accountability.

To summarize: We have good reasons to expand the scope of direct moral accountability to go beyond actions and include attitudes, character, and self, because (i) it is hard to defend a wholesale rejection of the wide range of blaming practices about things other than actions; (ii) attitudes, character, and self can indicate substandard qualities of will that some argue are sufficient for falling in the domain of moral accountability; and (iii) attitudes, character, and self can be so tightly connected to actions that drawing a sharp boundary of direct accountability in between is hardly plausible.

I think that this *expansive view on moral accountability*, the view that we are sometimes directly accountable for things other than actions, including attitudes, character, or self, is justified. But it is fair to say this remains controversial. Still, let us take the claim seriously for the moment, and consider the further question: *If* we are directly morally accountable for attitudes, character, and self, what does this entail about reactive attitudes? The first upshot is that the *Action-Based Objection* to inclusivism is no longer sound, since it is not the case that we are only ever directly morally accountable for actions. Instead, we can be directly morally accountable for attitudes, character, and self. All those are typically more self-focused than action-focused. This is obvious in the case of character and self, but also

rings true in the case of attitudes. Shame about an attitude is typically (though perhaps not always) connected with the thought that the attitude is likely to indicate something negative about one's character, self, or commitments; for instance, shame about one's racist attitude is typically connected with a thought that one is likely to lack a strong enough commitment against racism. The second upshot is that the *Control-Based Objection* to inclusivism is also going to be undermined. This is because attitudes, character, and the self are rarely within our voluntary control. Indeed, some defenders of the expansive view (e.g. Hieronymi, Sher) explicitly reject that any control condition is necessary for moral blameworthiness. The thought is that the kind of indication of flaws in quality of will is sufficient for moral blameworthiness, regardless of control. But another possibility is that accountability still requires control, but in a sense weaker than voluntary control. This may be a kind of 'rational control' (Smith 2005: 265) defined by whether something depends on one's rational judgments, or 'reasons-responsiveness' control defined by whether one or one's mechanism has a disposition to recognize and act according to the relevant reasons (e.g. Fischer & Ravizza 1998). Either way, the requirement of voluntary choice in moral accountability ends up being too demanding. We can then agree with the *Control-Based Objection* that the appropriateness condition of shame does not require voluntary control but maintain that a weaker control condition—probably that of rational control or reasons-responsiveness—still restricts the conditions under which shame, or a subcategory of shame, is appropriate. As an example, one may argue shame about an event is appropriate, in the sense that we have pro tanto reasons to feel shame, only when whether the event occurs depends on our rational judgments (in Smith's sense). So shame about a racist thought can

be appropriate on the assumption that whether one has that thought depends on one's rational judgments, whereas we do not even have reasons to feel shame about some of our bodily features, which we lack rational control over.

Further, I think the expansive view on moral accountability in fact provides more direct support for inclusivism. Consider the following, widely agreed connection between accountability and self-directed reactive attitudes:

If a person is negatively morally accountable for *X*, then it is appropriate for her to hold reactive attitudes toward herself with regards to *X*.

In cases where *X* is an action, it makes sense to take guilt to be the relevant reactive attitude. But if *X* is an attitude, character, or self, then I am inclined to think that shame may be better as a candidate emotion to fill in the conditional. Consider the three earlier cases again. It seems to me that the natural and paradigmatic emotion to feel for Amanda is shame about her sexist thought, for Barbara is shame about her cold character, and for Claire is shame about her cruel self. Guilt, by contrast, seems much less natural. How can Amanda feel guilt about a thought, Barbara feel guilt about a character, Claire feel guilt about her self, without their feeling guilt about any particular actions? This intuition vindicates the self-action contrast. Guilt is indeed primarily directed toward actions, and when it comes to attitudes, character, and self, the primary object of concern is more about how they reflect the person as *who she is*, which is better captured by shame rather than guilt.

The thought is that there is a useful division of labor between shame and guilt that we can draw. Both concern moral accountability, but guilt is more about the aspect in which one is accountable for particular actions, whereas shame is more about the aspect in which one's being accountable for things—attitudes, character, self, and perhaps some actions—reflects its root in some substandard features of the agent's self. Both facets are crucial for practices involving moral accountability. Though it may seem that moral accountability is always indexed to a particular substandard action, an actual or potential substandard self quite often falls in the scope of moral accountability as well. In fact, one common way for an individual to downplay a charge of blame is exactly to acknowledge her fault in the particular action, but not to in any way acknowledge the fault in her attitudes, character, or self. The above picture is desirable because it draws a clear division of labor between guilt and shame such that the two emotions play distinct yet interacting roles in our responsibility practices.

To recap, the argument in the current section can be formulated as follows:

The Expansion Argument for Inclusivism

- (1) The expansive view on moral accountability is correct.
- (2) If the expansive view on moral accountability is correct, then we have good reasons to endorse inclusivism.
- (3) Thus, we have good reasons to endorse inclusivism.

One may respond by either rejecting (1) or rejecting (2), and I will consider the main objections along those lines in section 5. It is worth noting that the *Expansion Argument* goes

hand in hand with the *Intertwinedness Argument*. The intertwinedness argument, as stated in the last section, is only convincing when we are shown that the familiar ways in which shame and guilt are differentiated fail to justify categorizing shame and guilt differently. This is what the expansion argument partly aims to show: the most familiar differentiations between shame and guilt, in terms of either the self-action contrast or voluntary control, would not justify different categorizations between shame and guilt, but rather show that shame and guilt are in charge of different aspects of our accountability practices. The intertwinedness argument is thus a convincing move with this supplementation, since, unlike in the case of resentment and fury, shame and guilt do not seem to be differentiated in ways that can justify different categorizations of the two emotions. From a different perspective, we can also view the intertwinedness argument as supplementing the expansion argument. One rejoinder people may have against the expansion argument is that, even though the most familiar ways of distinguishing shame from guilt do not contradict inclusivism, we may still speculate that there are other differentiations waiting to be found. However, the fact that shame and guilt are so deeply intertwined with each other renders this empirical speculation rather unlikely. More generally, we can combine the *Expansion Argument* and *Intertwinedness Argument* into a dilemma. On one hand, if one thinks there is no useful way of differentiating shame from guilt, then one already has no good resources to endorse exclusivism or divisionism, since one should already be convinced by the *Intertwinedness Argument*. On the other hand, if one wants to appeal to the familiar ways of differentiating shame from guilt to reject inclusivism, then one should be convinced by the *Expansion Argument* that those differentiations fail to justify categorizing shame and guilt

differently. Either way, there are good reasons to endorse inclusivism, regardless of one's substantial view on the nature of shame and guilt.

4. The Significance of Moral Accountability

As mentioned earlier, some may still resist inclusivism on the basis of some fairly substantial conceptions of moral accountability. They may suggest that the normative significance of moral accountability comes from its relation to, for example, *desert* or *demands and expectations*, and only guilt rather than shame captures this significance. However, I am inclined to think that inclusivism is fully consistent with these conceptions of accountability. In fact, the earlier arguments may bring these conceptions under a new light by locating the division of conceptual labor between guilt and shame in illustrative ways.

Let's start with desert. It is widely believed that moral accountability entails desert of sanction or reward. Pereboom refers to this as 'basic desert', in the sense that 'the agent would deserve to be blamed or praised just because she has performed the action, given an understanding of its moral status, and not, for example, merely by virtue of consequentialist or contractualist considerations' (2014: 2). It may be suggested that, because shame is not concerned with desert, it should not be an accountability-tied attitude. But why isn't shame concerned with desert? One answer is to again appeal to voluntary control (see Carlsson 2019): we only deserve adverse treatments for what we have voluntary control of, but shame about *X* can be appropriate even when we cannot voluntarily control *X*. But this is just the *Control-Based Objection* to inclusivism discussed earlier and faces the same difficulties.

Importantly, the inclusivist can reasonably maintain that accountability entails desert but deny that desert requires voluntary control. That is, we can pair both accountability and desert with weaker senses of control, like rational control or reasons-responsiveness. The desert-based objection then loses its force. The common examples against using control to restrict appropriate shame are usually those involving the lack of voluntary control, such as an attitude or a character that we cannot just choose or decide to alter. But it remains plausible that shame, or a subcategory of shame, cannot be appropriate if its object is in no way connected to our rational judgments or reason sensitivity at all. On this basis, we can maintain that shame, or a subcategory of shame, is in fact concerned with desert. It is also worth emphasizing that the deserved treatment should not be identified with harsh punishment. Though it is hardly plausible that Amanda, Barbara, and Claire deserve harsh punishment for their attitudes, character, and self, it is plausible that they deserve some adverse consequences, like other people's resentment, indignation, and the painful feelings of shame—of course, only proportional to how bad their attitudes, character, and self are.

Now turn to demands and expectations. One may think that the significance of moral accountability comes from those interpersonal demands or expectations that we hold people account to. Wallace defends exclusivism on the basis of this thought. Wallace appears to agree that shame is self-focused; he writes that when feeling shame one typically 'sees oneself as being all of a piece' and 'thinks of oneself as being thoroughly degraded' (1994: 241). But Wallace then suggests that 'if one violates a demand that one holds oneself to, it will be very hard to think of oneself as all of a piece' (1994: 241). Assuming that

reactive attitudes are instances of holding oneself to a demand, Wallace concludes that shame does not count as a reactive attitude.

However, it is unclear why there cannot be proper demands and expectations about the self-centered aspects of a person. On one way of understanding cases like Amanda, Barbara, and Claire, they exemplify that we do in our blaming practices hold other people to account for what we demand or expect of their attitudes, character, and self. And the most natural ways in which we view ourselves in relation to these demands or expectations are in the form of shame rather than guilt. To say the least, it is unclear why the function of shame cannot be to bring the agent herself to see the force of these self-focused reciprocal demands or expectations. Watson (2014) has made a similar objection to Wallace. As Watson points out, even if there were ‘pure shame cultures’, people in those cultures would still ‘be prone to a vivid sense of mutual expectations’ and ‘recognize obligations to one another’ (2014: 29). It is then reasonable to infer that, even in ‘our’ culture, shame or an important subcategory of it can involve the kind of demands and expectations that is central to accountability practices. One may then ask Watson why we in fact have two distinct sets of emotional attitudes—guilt and shame—if they are both concerned with our interpersonal demands and expectations. My earlier arguments in this paper help fill in this part of the story, since they show that there is still a useful division of labor between guilt and shame. Guilt involves holding oneself to more action-focused demands or expectations, whereas shame involves holding oneself to more self-focused demands or expectations, including those manifested in attitudes, character, and self.

I have considered two kinds of substantial conceptions of accountability—in terms of desert and interpersonal demands—and examined the plausibility of inclusivism in relation to these conceptions. These do not exhaust all possible conceptions, but the two examples provide important insights: First, theorists endorsing these conceptions frequently assume that accountability is only about actions or only about things within our voluntary control. Once these assumptions are rejected, it becomes clear that these conceptions can be consistent with, or even lend support to, inclusivism. Second, my earlier arguments for inclusivism can help illustrate how shame and guilt are concerned with different aspects of our accountability practices under these substantial conceptions of accountability.

5. Objections and Replies

5.1. Challenging the Expansive View

The expansive view on direct moral accountability remains controversial. An overall evaluation of this view, in contrast to the more restrictive view on the scope of direct moral accountability, can only be done by examining their various theoretical advantages and disadvantages and is beyond the scope of the current paper. In this section, I offer a response to those who may reject the expansive view. The purpose is to show that even the more restrictive view on moral accountability may nonetheless lead to reasons in support of inclusivism similar to the ones presented earlier.

To start with, note that, though the thesis that we are directly accountable for attitudes, character, and the self is fairly controversial, most theorists would agree that we can at least sometimes be *indirectly morally accountable* for an attitude, character, or self when and because (i) we are directly accountable for an action or a series of actions, and (ii) the formation of the attitude, character, or self is a reasonably foreseeable consequence of that action or that series of actions (cf., Fischer & Ravizza 1998). But then we could make a similar inference from the expansive view on the scope of indirect moral accountability to inclusivism, by examining those cases where we *lack epistemic access* to the earlier actions that we are directly accountable for. Consider a case as follows:

Lazy Professor. Dan is a mathematics professor. When he teaches Math 101, he consciously holds the belief that, in his class, male students are much more talented than female students in mathematics. But he neither expresses this sexist belief nor performs any discriminatory action based on this sexist belief. Further, Dan cannot recall what he did or decided not to do in the past that may have caused him more inclined to hold sexist beliefs. Still, it is reasonable to believe that Dan did something, such as skipping ethical trainings, that made him more likely to hold sexist beliefs.

What's special about this case is that the agent cannot really point to the thing that he was directly accountable for at an earlier time, though—let's stipulate—we have enough reasons to believe that there is such a thing that made him indirectly accountable for its

downstream consequence. Now consider the earlier biconditional between moral accountability and reactive attitudes again:

If a person is negatively morally accountable for X , then it is appropriate for her to hold reactive attitudes toward herself with regards to X .

What are the appropriate negative reactive attitudes for Dan to feel? Perhaps we can say it is appropriate for Dan to feel guilt about what he did earlier, but this does not seem to fully capture his correct emotional response. Imagine Dan apologizes to the students by saying ‘I feel really guilty about what I did earlier—which, I am sorry, I cannot remember what that was...’. It seems that such responses are far from satisfying, in the sense that they do not meet what we usually expect from agents in those situations. There should be something more. I contend that the ‘something more’ here can be captured in their emotional response toward the consequence of their directly accountable actions. So Dan should feel some negative reactive attitude about his sexist belief, even if he could not remember what he did earlier.⁷ We can then argue that, similar to the reasoning in section 3, shame rather than guilt is the more natural and paradigmatic emotion for Dan to feel about his sexist attitude, and this provides a good reason to endorse inclusivism.

⁷ Does this mean there is an asymmetry between guilt and shame, since guilt is appropriate only when one is *directly* morally accountable for an action? I am inclined to answer no. Suppose that Jim voluntarily decides to drink and drive but then involuntarily kills a pedestrian. If appropriate guilt is only about directly accountable actions, then it follows that Jim’s proper guilt should be only about drinking and driving. But it seems to me that he should also feel guilty about the action of hitting the pedestrian—an action that he is only indirectly accountable for. Therefore, I tend to think we should reject the assumption that appropriate guilt only concerns directly accountable actions.

More generally, the link between reactive attitudes and moral responsibility should not be restricted to direct responsibility. As an analogy, consider the view that we are only directly morally responsible for choices and decisions, whereas responsibility for outward behaviors is a more derived notion. Those who believe this view surely should not think that the only reactive attitudes are the appropriate emotions toward choices and decisions; they will include the appropriate emotions about outward behaviors too. The upshot is that, *if* we are indirectly accountable for attitudes, then we still have good reasons to include shame as an accountability-tied reactive attitude. And the antecedent here is widely accepted even among those who reject the expansive view on *direct* moral accountability.⁸

5.2. Challenging the Inference from the Expansive View to Inclusivism

What if one takes for granted the expansive view on direct moral accountability, but nonetheless rejects inclusivism? That is, why cannot one maintain that, in contrast to what I have suggested, we can feel guilty toward attitudes, character, and self? Why not think attitudes, or the processes involving forming attitudes, are themselves a kind of ‘mental action’?

⁸ However, it is worth noting that, on the picture discussed in the current subsection, shame is admittedly a less central reactive attitude than guilt, and its status as a reactive attitude is derived. This is not a concession that inclusivists would need to make if the expansive view on direct moral accountability is correct.

Smith (2011), for example, claims that guilty attitudes are in fact fairly common phenomenon. She writes that,

Perhaps we have caught ourselves taking secret pleasure in a close friend's uncharacteristic failure, or feeling distrustful of a loved one's fidelity, or viewing a stranger through the lens of an odious stereotype. Perhaps we have received the confidences of another with contempt, or have felt resentment rather than gratitude toward someone who has done us a kindness. It is quite common, in such cases, for people to say that they feel 'guilty' about these thoughts and attitudes, even when they are quite sure that they have not been, and will not be, expressed or acted upon in any way. The feeling of guilt, it seems, attaches to the mere having of these thoughts and attitudes, and is not inhibited by the knowledge that these mental states are and will remain wholly private. (2011: 235)

My reply is as follows. Even if we do sometimes feel guilty about these thoughts, this still does not justify treating guilt and shame differently when it comes to categorizing them as reactive or non-reactive attitudes. Admittedly, the observation does pose a serious challenge to my claim that shame is the more natural and paradigmatic emotion to feel when it comes to attitudes, character, and the self; however, it still does not accomplish what one needs in order to argue against inclusivism, namely, a way of differentiation between shame and guilt that justifies treating only the latter as an accountability-tied reactive attitude. If anything, the observation almost strengthens the thought underlying the *Intertwinedness Argument*.

Surely, many also feel shame about the above attitudes—taking secret pleasure, distrusting a loved one, stereotyping strangers, etc., and, when it comes to attitudes, guilt and shame appear even more intertwined than in the case of actions. It is even more common than in actions that guilt and shame overlap, are hard to differentiate, and frequently convert to each other. We thereby still have not been shown that the intertwinedness can be unraveled to support exclusivism or divisionism.⁹

5.3. Challenging from the Dark Side of Shame

A final line of objection to my arguments involves appealing to certain allegedly unique, negative features of shame. For example, one may worry that shame is not an ‘autonomous’ moral response, in the sense that it reflects not one’s own values but rather the values of other people in one’s society. The idea is that shame, even its ‘moral shame’ or ‘agential shame’ subcategory, seems to be strongly connected with one’s social position, such as the social power and oppression that one is subject to. This would be an argument against inclusivism if moral accountability does not involve such social factors. However, there has been promising developments of the view that moral accountability has a crucial social

⁹ We should also be cautious about drawing implications from those feelings we call ‘guilty thoughts’. I suspect that guilty thoughts in fact share more similarities with paradigmatic shame attitudes than with paradigmatic guilt attitudes. See, e.g. Goffin & Cova (2019), who conduct experiments to show that guilty pleasures are ‘not so much about guilt (what people do) but about shame and embarrassment (who they are)’ (1151).

dimension (see, e.g. Strawson 1962; McKenna 2018; Vargas 2018), which shame may be especially well-positioned to capture.

A different but related concern is that the moral relevance of shame is old news. A lot have been written on the positive moral value of shame, despite its apparently negative features.¹⁰ However, the focus there is usually on the moral relevance of shame in general, rather than about the particular connection between shame and moral responsibility. Even when some theorists (e.g. Manion 2002; Deonna, Rodogno, & Teroni 2012) do discuss this particular connection, they do not address some of the more recent developments in the moral responsibility literature, including the distinction between attributability and accountability, the expansive view on moral accountability, and the relation between direct and indirect accountability. My view is that we can establish a firmer defense of inclusivism by examining these specific choice-points in a theory of moral responsibility. Another way in which my arguments differ from the more general literature on the positive moral value of shame is this. My primary goal is to defend a connection between shame and moral accountability. I am inclined to think that defending this connection is in addition to, or even independent of, defending the positive moral value of shame. After all, there is the possibility that our practices involving holding people morally accountable do unfortunately have certain negative features and require significant revisions. We should not at the

¹⁰ See, for example, Velleman (2001), Calhoun (2004), Nussbaum (2004), Deonna, Rodogno, & Teroni (2012), Deonna & Teroni (2011), Manion (2002), Mason (2010), Maibom (2010), Thomason (2015, 2018).

beginning of our theorizing simply assume that whatever emotions central to our responsibility practices must be conducive to our being better moral agents, or that such practices are flawless, or that they cannot be modified or improved (see Vargas 2013: 76–7). This possibility can be easily overlooked when we do not disentangle the goal of defending inclusivism on one hand, and that of defending the general moral value of shame on the other.

6. Conclusion

I have presented novel support for including shame as a reactive attitude in the sense that concerns moral accountability, appealing to both the intertwinedness between guilt and shame, and the expansive view on the scope of moral accountability. This leads to various open questions. Where exactly should the scope of reactive attitudes be drawn? How about other self-directed and other-directed emotions such as pride, regret, and disgust? Does inclusivism about shame score better or worse once we consider these more general issues? I think they all lead to promising future projects, and I expect that the two aspects that I have focused on in the context of shame—the intertwinedness and interactions between different emotions, and the scope of moral accountability—will remain crucial for answering these broader questions.

Acknowledgement

For valuable comments and feedback on earlier versions of this paper, I want to thank many people including David Brink, Dana Kay Nelkin, Manuel Vargas, Douglas Portmore, Andreas Brekke Carlsson, Angela M. Smith, Monique Wonderly, Daniel Weltman, Richard Arneson, Rosalind Chaplin, Leonardo Moauro, Cory Davia, the audience at the CUNY Philosophy of Emotions Workshop in May 2019, the audience at the Moral and Political Philosophy Seminar at UC San Diego in October 2019. Special thanks to the two anonymous referees for this journal, whose comments have improved the paper significantly. Financial support for this work was provided by the LATAM Free Will, Agency, and Responsibility Project.

References

- Adams, R. M. (1985) 'Involuntary Sins', *The Philosophical Review*, 94/1: 3–31.
- Calhoun, C. (2004) 'An Apology for Moral Shame', *Journal of Political Philosophy*, 12/2: 127–46.
- Carlsson, A. B. (2019) 'Shame and Attributability', in D. Shoemaker (ed.) *Oxford Studies in Agency and Responsibility*, Vol. 6, 112–39. Oxford University Press.
- Clarke, R. (2013) 'Some Theses on Desert', *Philosophical Explorations*, 16/2: 153–64.
- Clarke, R. (2016) 'Moral Responsibility, Guilt, and Retributivism', *The Journal of Ethics*, 20/1–3: 121–37.

- Darwall, S. (2006) *The Second-Person Standpoint: Morality, Respect, and Accountability*. Harvard University Press.
- Deonna, J. A., Rodogno, R. and Teroni, F. (2012) *In Defense of Shame: The Faces of an Emotion*. Oxford: Oxford University Press.
- Deonna, J. and Teroni, F. (2011) 'Is Shame a Social Emotion?', in A. Konzelmann Ziv, K. Lehrer and H. B. Schmid (eds.) *Self-Evaluation: Affective and Social Grounds of Intentionality*, Vol. 116, 193–212. Dordrecht: Springer.
- Fischer, J. M. and Ravizza, M. (1998) *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge University Press.
- Frankfurt, H. (1971) 'Freedom of the Will and the Concept of a Person', *Journal of Philosophy*, 68: 5–20.
- Goffin, K. and Cova, F. (2019) 'An Empirical Investigation of Guilty Pleasures', *Philosophical Psychology*, 32/7: 1129–55.
- Graham, P. A. (2014) 'A Sketch of a Theory of Moral Blameworthiness', *Philosophy and Phenomenological Research*, 88/2: 388–409.
- Hieronymi, P. (2008) 'Responsibility for Believing', *Synthese*, 161/3: 357–73.
- Holroyd, J. (2012) 'Responsibility for Implicit Bias', *Journal of Social Philosophy*, 43/3: 274–306.
- Keltner, D. and Buswell, B. N. (1996) 'Evidence for the Distinctness of Embarrassment, Shame, and Guilt: A Study of Recalled Antecedents and Facial Expressions of Emotion', *Cognition and Emotion*, 10: 155–171.

- Lewis, H. B. (1971) *Shame and Guilt in Neurosis*. New York, NY: International Universities Press.
- Lewis, M. (2008) 'Self-Conscious Emotions: Embarrassment, Pride, Shame, and Guilt', in M. Lewis, J. M. Haviland-Jones and L. F. Barrett (eds.) *Handbook of Emotions*, 3rd ed., 742–56. New York: Guilford Press.
- Maibom, H. L. (2010) 'The Descent of Shame', *Philosophy and Phenomenological Research*, 80/3: 566–94.
- Manion, J. C. (2002) 'The Moral Relevance of Shame', *American Philosophical Quarterly*, 39/1: 73–90.
- Mason, M. (2010) 'On Shamelessness', *Philosophical Papers*, 39/3: 401–25.
- McKenna, M. (2018) 'Power, Social Inequities, and the Conversational Theory of Moral Responsibility', in K. Hutchison, C. Mackenzie and M. Oshana (eds.) *Social Dimensions of Moral Responsibility*, 38–58. Oxford University Press.
- Nelkin, D. K. (2016) 'Accountability and Desert', *The Journal of Ethics*, 20: 173–89.
- Niedenthal, P. M., Tangney, J. P. and Gavanski, I. (1994) "If Only I Weren't" Versus "If Only I Hadn't": Distinguishing Shame and Guilt in Counterfactual Thinking', *Journal of Personality and Social Psychology*, 67/4: 585–95.
- Nussbaum, M. C. (2004) *Hiding from Humanity: Disgust, Shame, and the Law*. Princeton University Press.
- Pereboom, D. (2014) *Free Will, Agency, and Meaning in Life*. Oxford University Press.
- Portmore, D. W. (2019a) 'Control, Attitudes, and Accountability', in D. Shoemaker (ed.) *Oxford Studies in Agency and Responsibility*, Vol. 6, 7–32.

- Portmore, D.W. (2019b) 'Desert, Control, and Moral Responsibility', *Acta Anal*, 34: 407–26.
- Rawls, J. (1971) *A Theory of Justice*. Harvard University Press.
- Sher, G. (2006) *In Praise of Blame*. Oxford University Press.
- Shoemaker, D. (2015) *Responsibility from the Margins*. Oxford University Press.
- Shoemaker, D. (2017) 'Response-Dependent Responsibility; or, a Funny Thing Happened on the Way to Blame', *Philosophical Review*, 126/4: 481–527.
- Smith, A. M. (2005) 'Responsibility for Attitudes: Activity and Passivity in Mental Life', *Ethics*, 115: 236–71.
- Smith, A. M. (2011) 'Guilty Thoughts', in C. Bagnoli (ed.) *Morality and the Emotions*, 235–56. Oxford University Press.
- Smith, A. M. (2012) 'Attributability, Answerability, and Accountability: In Defense of a Unified Account', *Ethics*, 122/3: 575–89.
- Strawson, P. (1962) 'Freedom and Resentment', *Proceedings of the British Academy*, 48: 1–25.
- Tangney, J. P. (1992) 'Situational Determinants of Shame and Guilt in Young Adulthood', *Personality and Social Psychology Bulletin*, 18: 199–206.
- Tangney, J. P. (1995) 'Recent Empirical Advances in the Study of Shame and Guilt', *American Behavioral Scientist*, 38: 1132–45.
- Tangney, J. P. and Dearing, R. L. (2002) *Shame and Guilt*. New York, NY: The Guilford Press.
- Tangney, J. P., Stuewig, J. and Mashek, D. J. (2007) 'Moral Emotions and Moral Behavior', *Annual Review of Psychology*, 58/1: 345–72.

- Taylor, C. (1976) 'Responsibility for Self', in A. O. Rorty (ed) *The Identities of Persons*, 281–99. Berkeley: University of California Press.
- Thomason, K. K. (2015) 'Shame, Violence, and Morality', *Philosophy and Phenomenological Research*, 91/1: 1–24.
- Thomason, K. K. (2018) *Naked: The Dark Side of Shame and Moral Life*. Oxford University Press.
- Tracy, J. L. and Robins, R. W. (2006) 'Appraisal Antecedents of Shame and Guilt: Support for a Theoretical Model', *Personality and Social Psychology Bulletin*, 32/10: 1339–51.
- Vargas, M. (2013) *Building Better Beings: A Theory of Moral Responsibility*. Oxford University Press.
- Vargas, M. (2018) 'The Social Constitution of Agency and Responsibility', in K. Hutchison, C. Mackenzie and M. Oshana (eds.) *Social Dimensions of Moral Responsibility*, 110–36. Oxford University Press.
- Velleman, J. D. (2001) 'The Genesis of Shame', *Philosophy and Public Affairs*, 30/1: 27–52.
- Wallace, R. J. (1994) *Responsibility and the Moral Sentiments*. Cambridge, MA: Harvard University Press.
- Watson, G. (1996) 'Two Faces of Responsibility', *Philosophical Topics*, 24/2: 227–48.
- Watson, G. (2014) 'Peter Strawson on Responsibility and Sociality', in D. Shoemaker (ed.) *Oxford Studies in Agency and Responsibility*, Vol. 2, 15–32.
- Westlund, A. C. (2003) 'Selflessness and Responsibility for Self: Is Deference Compatible with Autonomy?', *The Philosophical Review*, 112/4: 483–523.

University of California, San Diego, USA