

IX—EQUAL OPPORTUNITY: A UNIFYING FRAMEWORK FOR MORAL, AESTHETIC, AND EPISTEMIC RESPONSIBILITY

DANA KAY NELKIN

On the one hand, there seem to be compelling parallels to moral responsibility, blameworthiness, and praiseworthiness in domains other than the moral. For example, we often praise people for their aesthetic and epistemic achievements and blame them for their failures. On the other hand, it has been argued that there is something special about the moral domain, so that at least one robust kind of responsibility can only be found there. In this paper, I argue that we can adopt a unifying framework for locating responsible agency across domains, thereby capturing and explaining more of our actual practices. The key, I argue, is to identify the right conditions for being morally accountable, which I take to be a matter of having an opportunity of a good enough quality to act well. With this account in hand, I argue that we can adopt a unifying framework that allows us to recognize parallels across domains, even as it points the way to important differences among them.

I

We naturally speak about moral obligations ('you ought to have kept your promise'), and we speak about epistemic ones and even sometimes what look like aesthetic ones too ('you ought to have known that the polls were within the margin of error', 'you should have done better with that painting'). Similarly, we blame and praise people for epistemic and aesthetic transgressions and achievements, as well as moral ones. Consider a couple of examples. Although in the aesthetic realm we often focus on the praiseworthy (in an interesting contrast to the moral, where the focus tends to be the blameworthy), there are cases in which we blame artists for bad pieces of art, such as a recently replaced and widely condemned statue of

Lucille Ball in her hometown in upstate New York.¹ Similarly, consider the case of the paintings of Margaret Keane, once attributed to her husband, Walter Keane.² In a now famous review of *Tomorrow Forever*, an enormous painting of one hundred big-eyed children commissioned for the 1964 World's Fair, the *New York Times* critic, John Canady, wrote,

Mr Keane is the painter who enjoys international celebration for grinding out formula pictures of wide-eyed children of such appalling sentimentality that his product has become synonymous among critics with the very definition of tasteless hack work. (Canady 1964)

This is clearly not simply criticism of the painting. And it seems to me that it is not merely an expression of a judgement of Mr Keane's lack of talent or taste or aesthetic virtue. The idea that Mr Keane is accused of having 'ground out' the pictures and that the process of creation is described as 'hack work' seems to suggest his being held to account for the failure to meet a standard that he could have been expected to meet.

And in the epistemic realm, we find many examples of what appears to be recognition of both blame and obligation. For example, we say things like 'Think!' and 'You can surely put the pieces together!' and 'You ought to have figured that one out.' Cases of blame toward the self for epistemic transgressions can be particularly vivid. For example, consider a case in which one makes overly optimistic inferences from poll results on the eve of an election, having failed to take the margin of error properly into account. Based on personal and anecdotal evidence, it seems that one can feel a particularly strong kind of self-blame about having committed that sort of error.³

For these reasons, it is natural to conclude that our moral, aesthetic, and epistemic practices should be treated in highly

¹ See Stack (2016) for the story and photographs.

² Their story is documented in a recent Tim Burton film, *Big Eyes* (2014). The fact that Walter Keane took credit for Margaret Keane's work adds a serious moral dimension to the story, a point to which I will return. But here I focus on the critical reception of the work, and the public condemnation of Walter Keane, who at the time was believed to be its creator. See Ronson (2014) for details of the story and photographs, including one of Margaret Keane with *Tomorrow Forever*.

³ It is notable that several interesting recent treatments of epistemic responsibility have either taken it as a presupposition that there is such a thing as epistemic blame (for example, Meylan 2019 and Rettler 2018) or aimed to explicate accounts of epistemic blame (see Tollefsen 2017 and Brown 2020).

parallel ways, at least when it comes to the realm of holding people responsible, praiseworthy, and blameworthy. Importantly, though this paper focuses on these three realms, I believe that much of what I have to say about the parallels among them can extend to other realms as well, such as the athletic. For present purposes, the aesthetic and epistemic, when compared to the moral, provide an already helpful diversity of features for exploration.

It is also worth noting that although moral responsibility tends to be the focus in influential discussions of free will and responsibility, it is often noted in that context that moral responsibility is only one type of responsibility. To take just one example, Robert Kane, in *The Significance of Free Will* (1996), points out that there are at least ten things that are of great importance and that seem to presuppose that we are free agents. These include moral responsibility, but also desert for one's achievements and accomplishments, specifically including great works of art and symphonies for which their creators are responsible.

At the same time, there are clear asymmetries between the moral and the epistemic, and also between the moral and the aesthetic, which might seem to doom any hope for a genuinely parallel treatment. For example, as many have pointed out, unlike actions or omissions, which are clear objects of moral obligations, praise, and blame, belief does not seem to be the kind of thing over which one has control (see, for example, Alston 1988). When it comes to the aesthetic case, many have doubted that we have obligations in the way that we do in the other cases, and the idea of aesthetic blame is relatively rarely discussed in the aesthetics literature.⁴ Further, moral blame has seemed to many to be governed by a number of interpersonal norms that don't seem to have parallels in *either* the pure epistemic or pure aesthetic case.

Despite these challenges, I believe that the prospects are promising for a unifying framework that applies in all three cases. In particular, I will defend the idea that we can be responsible, blameworthy, and praiseworthy in all three domains, and that in each case one's degree

⁴ But see Eaton (2008) endorsing the existence of aesthetic obligation and responding in part to the scepticism of Hampshire (1954), and Archer and Lauren (2017) for recent discussion of aesthetic supererogation.

of blameworthiness or praiseworthiness depends on the quality of one's opportunity.

There are different approaches one might take to the project. It is notable that the question of whether there is such a thing as epistemic responsibility—and if so, what it is—arises in a number of different contexts. It emerges from discussions of deontological views of justification in epistemology; in moral and legal contexts focused on culpable ignorance; in discussions of parallels between practical and theoretical reflection; and more. For this reason, there are many different ways to approach the topic: one can begin in a more 'bottom-up' way, starting with a particular case and considering specific principles that might govern it, or one can begin in a more 'top-down' way, seeing whether a general framework can accommodate a variety of cases. In this paper, my strategy will be closer to the latter.

In §II, I set out a framework for thinking about moral responsibility, and in §§III and IV I show how it can be applied to the cases of epistemic and aesthetic responsibility. In §§V, VI and VII, I consider objections, allowing for further elaboration of the framework.

II

Moral Responsibility. To begin, it is important to clarify the relevant notion, or notions, of moral responsibility. Following Watson (to start) we can distinguish between two notions of responsibility. According to Watson, one is responsible in the attributability sense if one's actions reflect one's having adopted an end, one's having committed oneself to a certain conception of value. To blame someone in this sense is to attribute a moral fault to an agent (Watson 2004, p. 266). In contrast, one is responsible in the accountability sense if it is appropriate to make moral demands of one as an agent. And our practices of holding one another accountable 'involve the imposition of demands on people' (Watson 2004, p. 273). In turn, according to Watson, imposing demands is a matter of laying it down that 'unless the agent so behaves she will be liable to certain adverse or unwelcome treatment' or sanctions (Watson 2004, p. 275). Sanctions are in turn connected with reactive attitudes such as indignation, resentment, and disapprobation in the following way: either they are themselves disagreeable when experienced by their

targets or they involve dispositions to treat others in generally unwelcome ways. Because the prospect of adverse treatment arises in this way, questions of fairness arise in connection with accountability, and it is for this reason that many have thought being able to do otherwise is essential for accountability. In order to be blameworthy in the accountability sense, the demands in question must be fair or just.⁵

It is responsibility in this second, accountability sense that is then taken to be central in classic debates about whether we can be morally responsible agents. (One important position takes it that being responsible in the attributability sense is sufficient for being responsible in the accountability sense, but this is a substantive position, to which we will return.⁶) And it is responsibility in the accountability sense that is taken by many to be at the heart of debates about the justification of moral emotions such as resentment and indignation, the making of moral demands, and blaming practices and punishment, and related to debates about free will.

What is the relationship of desert to responsibility in this sense? Intuitively, and as is sometimes simply taken for granted, when one is blameworthy in the accountability sense, when one has violated a legitimate moral demand, one is deserving of sanctioning responses.⁷ In addition, as I have argued elsewhere, being deserving of negatively valenced responses and being blameworthy in the accountability sense have the same satisfaction conditions. My view of what those conditions are is controversial, and while I do not have the space to offer either a full explication or defence here, I hope to say enough to make the account plausible.

⁵ Shoemaker (2015) argues for a tripartite distinction among attributability, answerability, and accountability. But his account of accountability resembles Watson's in at least some key ways (p. 87), and I will focus on what I take to be common ground here.

⁶ See, for example, Smith (2008, 2012). Smith (2015) argues that there is really just one notion associated with our practices, and calls it 'answerability'. While I believe that she is correct that some of the examples that have been put forward in support of the distinction are inadequate, there are other cases that she does not consider which seem to me to do a better job of illustrating the distinction, including children and people who are intoxicated through no fault of their own, for example. Further, I take it that in many cases we readily speak of blaming people for moral faults in a sense that does not entail that we hold them accountable for them, and fully excuse them. At the same time, I am sympathetic to the idea that using '*responsible*' in these contexts is a stretch of ordinary language in a way that using 'blameworthy' is not. (See Levy 2017 for a denial that attributability really captures a notion of responsibility.)

⁷ Recently, however, a number of theorists have explicitly questioned this mutual entailment. See, for example, McKenna (2012) and Shoemaker (2015), and see Nelkin (2016) for a response.

In thinking about accountability, understood fundamentally in terms of the aptness of demands, it seems that in order to be accountable one must have the ability to comply with the relevant demands. Given an understanding of the contents of the relevant demands as demands to act (or not act) in certain ways for the right reasons, it seems that the ability to so act (in the right ways, for the right reasons) is required in order for demands to be apt.⁸ In particular, one needs to have the opportunity to avoid wrongdoing. As I understand ‘opportunity’ here, it is a function of both the agent’s capacities and the congeniality of the situation. Further, opportunities can be better or worse on either or both of these dimensions, and blameworthiness and praiseworthiness in the accountability sense can track this scalar aspect of the quality of agents’ opportunities.

To see this more concretely, note that in order to be *blameworthy at all* in the accountability sense, one must have not only an opportunity to avoid wrongdoing, but also a *good enough* opportunity. In previous work, David Brink and I argue that the idea that blameworthiness depends on the quality of one’s opportunity to avoid wrongdoing best captures the wide variety of commonly recognized excuses that otherwise look very heterogeneous (Brink and Nelkin 2013). For example, it explains why we recognize excusing conditions ranging from ones that compromise our normative competence (via either cognitive or volitional impairments) to ones that impose situational constraints such as conditions constituting duress. What ultimately brings these together is that both such impairments and such constraints can in their own ways result in a lack of opportunity (or a lack of a high enough quality of opportunity) to avoid acting wrongly. Working backwards from excuse to blameworthiness in the accountability sense, we can see that the latter is then instantiated just when one acts wrongly and at the same time possesses an opportunity of high enough quality to avoid wrongdoing, or more positively, to do the right thing for the right reasons.⁹

These same conditions plausibly capture desert. Consider blameworthy action first. Intuitively, it seems that one is deserving of a

⁸ For further elaboration of this sort of view, see Nelkin (2011, 2015). Wolf (1990) does not use the language of ‘accountability’ or understand responsibility in part in terms of demands, but she offers a similar set of satisfaction conditions for what she there calls ‘free and responsible’ action.

⁹ Though I favour an interpretation of ‘opportunity’ on which having an unexercised opportunity is compatible with determinism, for current purposes I believe it is best to be ecumenical and not take a stand in the debate between compatibilists and incompatibilists.

sanction, or negative response, only if one has an opportunity to do the (or a) right thing for the right reason. More specifically, one is deserving of a sanction only if one has a reasonably high-quality opportunity to do the (or a) right thing for the right reasons and failed to take it. In fact, it is plausible that the conditional goes both ways. Having a fair opportunity to do the right thing for the right reasons and failing to take it is not only necessary but sufficient for being deserving of sanction.

It is important to note that nothing follows without additional premisses about the goodness of receiving a sanction or about its permissibility or fairness, all things considered. Being deserving of a sanction does not make it good for one to receive that sanction, nor does it even by itself provide others with a reason to sanction. But under certain circumstances it can be part of such a reason, even an overriding one, and in this way desert is indirectly linked with fairness. There is, then, an additional reason, aside from intuitiveness, to see that having the relevant fair opportunity to do the right thing for the right reasons (thereby avoiding wrongdoing) is necessary for desert by virtue of explaining the common connection between desert and fairness.

Intuitively, the idea of what one does with one's opportunities also fits with desert of benefit as well. When one does the right thing when it is really difficult, say, then one deserves particularly positive responses. Thus, what I will hereafter refer to as the 'Quality of Opportunity' view captures the idea that acting badly with a sufficiently high quality of opportunity is necessary and sufficient for desert and blameworthiness in the accountability sense, and that doing well with a sufficiently low quality of opportunity (where factors include a high degree of difficulty or sacrifice) is necessary and sufficient for desert and praiseworthiness in the accountability sense.

Now, let us apply this to a particular case that will also allow for adding one piece to the picture. Consider Santana, who betrays a confidence of a close friend. A natural response is that her action is morally faulty, and thus blameworthy in the attributability sense. But beyond this, given the information we have so far, it also seems a defeasible but reasonable response to think that she is morally accountable and blameworthy in the accountability sense for her actions, and that she is deserving of a negative response for them.

But now suppose we find out that Santana's opportunity was of low quality; for example, that she was drugged so that her normal

self-editing disposition was masked in some way. Or suppose we find out that she was raised in a cult that cultivates self-centredness in its children, so that she couldn't see that her behaviour was out of bounds. This might lead us to suspend our judgement that Santana is accountable and blameworthy in the accountability sense. Arguably, she still acted badly in a way that reflected her own ends and values, and we can rightly take her to be blameworthy in the attributability sense, but it is much less clear that we would be right to hold her accountable, or that she is deserving of sanction in these cases.

The view just sketched can accommodate these thoughts. It requires us to see that acting responsibly is a matter of one's having opportunities—and in particular, the opportunity to act well.¹⁰ Understanding responsible action in these terms allows us to explain our tendency to excuse Santana in a principled way, as based on a lack of opportunity to act well in the circumstances.

Imagine, however, that Santana arranged to take the drug, knowing that it would have the effect of loosening whatever inhibitions would have prevented her from betraying her friend's confidence. In that case, it is natural to say that she is not thereby off the hook and is blameworthy despite not having an opportunity to do better at the time she reveals the confidence. This suggests that we need to add what is often called a 'tracing' component to the account. The idea is that to be blameworthy for a given action or omission, one must have a good enough quality of opportunity to avoid wrongdoing, or one's blameworthiness must trace back in the right way to an earlier time at which one met the conditions for responsibility. I return to this idea in §V.

In grounding accountability in the nature of one's opportunities to meet relevant demands (and the obligations or standards on which these rest), the Quality of Opportunity view is best categorized as a 'control' view, in contrast to views that emphasize the nature of evaluative judgement or quality of will expressed in action or attitude. As we shall see, this opens the view to particular challenges

¹⁰ This means that the view is asymmetrical, in the sense that it requires the ability to do otherwise for morally responsible action only if one acts badly. On some competing views, one must have the opportunity to do well *and* the opportunity to do badly, or at least otherwise, even if one acts well. (See, for example, Haji 2016.) I have defended the view in the text in more detail elsewhere (see Nelkin 2011). For now, it is important to note that the two kinds of views will converge in the case of blameworthy action in requiring an opportunity to act well.

that so-called ‘quality of will’ views do not face in attempting the unifying project. But in the next sections, I will try to show that the Quality of Opportunity framework has natural application in the aesthetic and epistemic realms, respectively.

III

Aesthetic Responsibility. As we have seen, there are many apparent parallels between the moral and aesthetic cases. We take people to be morally responsible for their actions, we blame and praise them, we take them to be deserving. The same goes for the aesthetic case: we take people to be responsible for their artistic achievements (and failures), we blame and praise them, we take them to be deserving. Are we then compelled to conclude that our practices surrounding artistic achievement and failure presuppose all that is required for accountability, including, on the framework just sketched, sufficiently high quality opportunities to act well?

A natural first response is that the framework just set out provides an elegant way of offering a ‘no’ answer. Artists might be praiseworthy and blameworthy in a sense of responsibility that parallels the moral attributability sense, and for this, as we have seen, they need not have had opportunities to do well. Artists can express great aesthetic virtues (or vices) in their art as aspects of their true selves, and on this basis be praiseworthy in an important way. Perhaps this captures all that we really need to explain and justify our practices, and does so in an admirably minimalist way. As Susan Wolf (2015) has recently argued, responsibility in the attributability sense is often wrongly seen as a ‘shallow’ kind of responsibility, when it is actually quite deep, reflecting the true self of the agent, person or artist. It allows us to praise and blame artists for their work, and in ways that depend on our taking that work to reflect their values and judgments and self-defining traits.¹¹ This response locates a fundamental asymmetry between moral and aesthetic responsibility while still recognizing a parallel at the level of attributability.

While I agree that responsibility in the attributability sense is indeed ‘deep’ and important, I do not think it captures all there is

¹¹ Interestingly, Wolf uses the case of the artist, along with cases in which people’s non-moral traits, such as having a sense of humour or charm, are attributable to them, in order to support the very idea that responsibility in the attributability sense is in fact deep.

when it comes to aesthetic responsibility. For it seems that we do hold others to account for bad art, and we take them to be deserving of blaming and praising responses that might be burdensome or beneficial in a way that seems parallel to the moral case. Moral reasons and norms are of course different in kind from aesthetic reasons and norms, but there is a parallel to be made out that can help explain why it is that we can be accountable for bad aesthetic choices and actions, as well as moral ones. In each case, one can meet, exceed, or fall below the relevant standards, and be called to account if one fails. I tentatively propose, then, that a great deal of the framework set out for moral responsibility can be transposed to the aesthetic case.¹²

To see how this would work, imagine that the creator of the ‘Bad Lucy’ sculpture or the painter of *Big Eyes* were capable of creating great art with relative ease, and just decided to cut some corners instead. My intuition in these cases is that blame could be perfectly appropriate. (In order to avoid difficult questions about when and from whom expressed blame is appropriate, we can simply consider whether it would be appropriate for the artists to blame themselves. And here blame seems perfectly apt.) But if, through no fault of their own, they simply could not do any better than they did, then, while we might criticize the product, and even attribute negative artistic traits to the artists, it would be hard to work up an aesthetic analogue to indignation. This suggests that opportunity is required for justified blame—at least blame that is associated with the failure to

¹² My disagreement with Wolf on this point might depend in part on a disagreement that is even more fundamental, regarding the relationship between attributability and accountability. Notably, Wolf also claims that attributability is completely distinct from accountability, so that it is not even a necessary condition for it. I disagree, for reasons that I believe can be traced to a different understanding of accountability. In introducing the notion of accountability, she writes that ‘holding someone accountable involves making him liable to blame and punishment’ (Wolf 2015, p. 134). But later, when arguing that attributability is not necessary for moral accountability, liability to *blame* drops out, and liability to penalty takes centre stage, making it hard to distinguish from legal liability. She writes, ‘If [the psychopath who is not responsible in the attributability sense] acts criminally, knowing that the acts are forbidden, and that he will be punished for them if he gets caught, and if, as I have stipulated, he can control his behavior, then it does not seem unfair to impose the penalty’ and she goes on to conclude that he can thereby be accountable (Wolf 2015, p. 138). In contrast, on my view, what is demanded of one in the moral case is not just certain actions or the avoidance of forbidden ones, but also doing things for the right reasons. This implicates the ends for which the agent acts, and requires that he be responsible in the attributability sense to be morally accountable. (See Nelkin and Rickless 2014 for a defence of the claim that intentions are relevant to moral obligations.)

meet applicable standards and demands that rest on the recognition that one could have reasonably been expected to have done better.

Other interesting questions arise about how much of the detail of the moral framework transposes to the aesthetic. The scalar aspect of the Quality of Opportunity account in the moral realm seems to have a counterpart in the aesthetic case, and the nature of the factors on which degrees of blameworthiness and praiseworthiness arguably have something in common as well. For example, it seems that difficulty of doing the morally right thing can enhance one's praiseworthiness when one succeeds, and mitigate one's blameworthiness when one fails. Is difficulty also a factor in the aesthetic case? In so far as responsibility for a piece of art is seen as an instance of the more general category of achievement, a case can be made that difficulty is a factor in the degree of praise or blame that is appropriate.¹³ If two artists created two equally beautiful paintings but one had to work in very difficult conditions, then, though they might be equally praiseworthy in the attributability sense, it seems that the one in the less conducive situation is more praiseworthy in the sense that is analogous to moral accountability. There is much more to say here about the way in which difficulty affects aesthetic praiseworthiness and blameworthiness; for now, I note that this point also applies in the moral case.

There is one aspect of the moral framework that, at least on the face of it, lacks a counterpart in the aesthetic realm. Failing to develop or exercise one's aesthetic (or other) talents does not automatically entail that one is blameworthy. In the moral case, one doesn't 'opt in' to the arena, whereas when one becomes an artist or takes on a particular aesthetic project, one in some way opts in. Once one has opted in, so to speak, and meets the conditions for responsible agency in that realm, one is then governed by a set of standards and becomes a candidate for a special sense of blame and praise.¹⁴

¹³ For interesting discussion of the nature of achievement as essentially something hard to do, see Bradford (2015).

¹⁴ Tønissen (MS) discusses the subtly different ways in which aesthetic obligations might be 'opt in'. It is, of course, a controversial question whether there are such obligations, or even objective standards. For current purposes, I am simply assuming that there are objective aesthetic standards, but I mean for the view to be as ecumenical as possible as to what they are, how context-dependent they might be, and so on, given that assumption. It is also the case that opting in to a set of practices might straddle different sets of standards. For example, if I take a position as a professional artist, then I have opted in to a set of aesthetic standards, but also perhaps to a particular set of community standards. Finally, it might be thought that there are *some* aesthetic standards that apply to all of us, and that do not need

Thus, aesthetic responsibility is not parallel in every respect to moral responsibility. But much of the framework of moral responsibility sketched earlier, including the distinction between attributability and accountability and the idea that opportunity is at the heart of the conditions for accountability, applies neatly in both cases.

Now let us turn to the epistemic case.

IV

Epistemic Responsibility. As we have seen, there appear to be many parallels between the moral cases and the epistemic cases. Are we then compelled to conclude that our practices surrounding epistemic achievement and failure ought to be treated in a similar way?

Just as with the aesthetic case, a natural first response is that the framework just set out provides an elegant way of offering a ‘no’ answer. We might be praiseworthy and blameworthy for epistemic successes and failures in a sense of responsibility that parallels the moral attributability sense. One’s epistemic states and processes reveal oneself to be a certain sort of epistemic agent—one who is attuned to the evidence; good at assimilating disparate information; prone to self-deception, and so on. Perhaps this captures all that we really need to explain and justify our practices. But, as before, I think this is not the case.

For it seems that we do hold others to account for epistemic failures, and we take them to be deserving of blaming and praising responses that might be burdensome or beneficial in a way that seems parallel to the moral case. At least some are willing to speak of epistemic obligations (and even more of norms).¹⁵ Further, we also make demands that seem to invoke epistemic norms: ‘Think!’,

opting into. See [Press \(1969\)](#) and [Lopes \(2018\)](#). Press has in mind obligations to cultivate taste and judgement rather than to create, which has been the paradigm case discussed here. [Lopes](#) in places suggests that there are standards that attach to us all, despite the fact that on his view, aesthetic reasons attach in the context of particular aesthetic practices; for example, he writes that even without developed aesthetic practices our distant ancestors were ‘provisioned with plenty of low-level aesthetic goodness . . . In view of the goodness of these things, they had aesthetic reasons to act in ways that turned out to enlarge their aesthetic inventory’ ([Lopes 2018](#), p. 112).

¹⁵ For interesting recent discussions of epistemic obligations on the positive side of our responsibility practices, see [Hedberg \(2014\)](#) for a defence of supererogation in the epistemic realm, and [McElwee \(2017\)](#) for discussions of supererogation in the epistemic realm, as well as the prudential and etiquette realms.

‘Put the pieces together!’, ‘Stop deluding yourself!’, ‘Try harder!’ And we blame in ways that go beyond the making, or even uttering, of evaluative judgements: ‘Can’t you *see*?!’ I tentatively propose, then, that a great deal of the framework set out for moral responsibility can be transposed to the epistemic case.

Just as with opportunities in the moral case, the quality of one’s epistemic opportunities will be a function of one’s situation, on the one hand, and one’s capacities, both cognitive and volitional, on the other. So, to return to the case of the reader of the poll results on the eve of an election, whether one is blameworthy for not realizing that the results were within the margin of error in the key states (say!) depends on a variety of factors, including the evidence available (were individual state polls reported or simply aggregated?), whether one had a solid understanding of the relevant concepts (does one know what ‘margin of error’ means?), whether one was subject to such strong wishful thinking that one could not recognize or control for it, and so on. Of course, it will be difficult in particular cases to know the relevant empirical facts about one’s capacities and situation at the time. But whether or not one is blameworthy depends on them.

Or consider a case in which someone gets an answer wrong on a maths test. Whether one is blameworthy for getting the wrong answer will depend on a number of factors, including one’s cognitive skills. If one simply doesn’t have any knowledge of calculus, say (and suppose this is through no fault of one’s own—more on which later), then it doesn’t seem that one would be blameworthy for one’s failure. Here too, then, it seems that much of the framework of moral responsibility sketched earlier applies. Other sorts of cases that might even more easily bring out a kind of analogue of indignation and other attitudes associated with blame are cases in which people appear to make clearly incorrect inferences about the causes of climate change or about dogmatic religious claims. On the one hand, these sorts of cases are quite powerful, and, on the other, they deserve further scrutiny; I return to them in the following sections, where I consider objections that will allow for elaboration of the account, particularly as it applies in the epistemic case.

V

A Challenge Concerning the Inescapability of Mixed Cases. Consideration of a case such as bad formation of beliefs about climate change might naturally lead to the objection that this is really just a moral case in disguise, and that rather than there being a parallel to moral responsibility in the epistemic domain, what we really have here is just moral responsibility once again. People have a moral obligation to arrive at beliefs that will not risk serious harm, and this is a moral obligation at its foundation. Similarly, it might be thought that in the Lucy case the artist was commissioned to create a likeness, and this clearly did not happen. So, there seems a breach of contract, and perhaps more than that, a failure to represent the subject in a good light—these are moral (and legal) failings. When we blame in light of them, we are engaged in familiar practices of moral responsibility.

These are clearly ‘mixed’ cases, in the sense that they both have moral implications, and the objection that flows from this observation needs to be addressed.¹⁶ Perhaps it could be sidestepped if we could find ‘pure’ cases—ones without any moral implications. The *Tomorrow Forever* case seems a better candidate for being ‘pure’. By all accounts, the World’s Fair got exactly what they contracted for: a picture in a style that was widely popular. No contract was breached, and fair organizers and fair-goers were pleased. This is not an obviously pure case, however, as it could be argued that the artist was in some way representing all purveyors of his craft, and failed to do well by his fellow artists. But the review by Canady does not mention these things, and seems to be focused on the artist’s having

¹⁶ Matheson and Milam (2020) provide a strong case that many instances of blame, including cases that might initially seem to be non-moral, such as cases of blaming athletes for sub-standard performances, are really moral after all. I am sympathetic to the idea that there are a great many cases that have at least some moral component, even if moral blame does not exhaust the kinds of blame involved. But I believe that we take the point about many cases having a moral component to have differential significance. As I read Matheson and Milam, this point plays a large role in supporting their claim to have shifted the burden to those who would see even a *pro tanto* reason to blame in non-moral ways. Since many take blameworthiness itself to provide a *pro tanto* reason for blame, I take it that they ultimately see the burden falling on those who would recognize parallels for non-moral blameworthiness of the kind I have defended here. While a full assessment of burden possession would require more space than I have here, I take it that being able to provide some good candidates for ‘pure’ non-moral cases as I do in the text neutralizes this burden. And, more importantly, being able to provide an appealing unified account, together with its flexibility on a variety of dimensions to be discussed, either discharges such a burden or shows that the burden lies elsewhere.

come up short with respect to aesthetic demands. So it might be difficult to be sure that we have located pure cases.

Turn to the epistemic realm. Some cases are clearly mixed. For example, there are cases such as that of the doctor who looks (or fails to look) at the patient's chart to figure out whether there is a risk of fatal allergic reaction to a particular medication that she is considering prescribing. This is a mixed case, in that there seems to be a moral obligation for the doctor to perform in certain ways, as well as epistemic norms governing her gathering and analysis of evidence as well as her inferences. In contrast, consider cases that are better candidates for pure cases: I make a wrong inference based on polling results, or I am engaged in a mathematical calculation, or in committing obscure dictionary definitions to memory. Nothing moral hangs on my success or failure—we can suppose that no one else's well-being will be affected positively or negatively either way, no one's rights will be violated, and so on. But it seems that I might still be subject to *epistemic* (non-moral) norms, and that I can be blameworthy or praiseworthy in an epistemic sense for my performance. This might take a non-trivial form, and at least one case in which this arises is in the case of self-blame.

Still, it might be argued that all cases are mixed once we examine them thoroughly, and that therefore all responsibility is ultimately moral (or, more broadly, practical). W. K. Clifford ([1876] 1999), for example, famously wrote, 'It is wrong always, everywhere, and for anyone, to believe anything upon insufficient evidence'. In the context, the 'wrongness' at issue is moral in nature. This view, sometimes called 'moral evidentialism' (Ryan 2015), suggests that there are no non-moral epistemic cases. But even if Clifford is correct, and moral norms apply in all cases of belief formation, we can still conceptually separate moral and epistemic norms, and blame or praise accordingly for how one fares with respect to them. The moral norm on this view might be simply to believe on evidence that is sufficient to epistemically justify, or make epistemically rational, one's beliefs. In that case, the epistemic norms are themselves purely epistemic, and in addition, there is a moral norm that enjoins us to respect certain epistemic norms. We can then distinguish between one's epistemic performance and one's moral performance, even if one is always doing both sorts of things when one acquires beliefs.

But moral evidentialism does not provide the only challenge to the idea that we can distinguish different kinds of cases as our starting

point for offering parallel treatments. A related challenge is presented by recent discussions of pragmatic encroachment, the idea that practical factors are relevant to determining whether or not a belief meets the standards for justification, knowledge, or rational belief. Moral encroachment is a more specific version of this idea, namely, that moral factors are relevant in determining whether or not a belief meets the relevant epistemic standard (see Guerrero 2007, Stanley 2005, Stroud 2006, Weatherson 2008, and Moss 2018 for discussion). For example, suppose that Mia and Sky are driving to deposit a cheque at the bank, and they remember that last year Saturday hours were 12–4 p.m. This might be all that is needed for their belief that the bank will be open when they arrive at 1 p.m. to be justified if they do not need the funds until the following week. But if they need the funds to cover a medical procedure the next day, then the very same evidence that otherwise would have served to justify their belief does not do so in this case with higher stakes. A paradigm case of moral encroachment is one in which a person has evidence that her friend has done something horrible, evidence that would be sufficient to believe that a stranger about whom she had all the same information had done it. But given that in this case the evidence points to her friend, she is not justified in believing it, and in fact should not (see Stroud 2006). In both of these cases, factors other than features of the evidence determine whether a belief is justified or rational and whether one ought to hold it.

Now, it is true that on some views that recognize pragmatic encroachment a neat conceptual division between epistemic and practical or moral norms would not be fully possible, in so far as, according to these views, how much evidence one needs for justified belief or knowledge varies depending on the stakes. This means that we cannot offer the same answer as the one above in response to moral evidentialism by pointing to a conceptual distinction between different kinds of norms. And if we cannot distinguish between different kinds of norms, then the implicit idea of there being different normative realms in the first place for which we can recognize parallel responsibility frameworks will be undermined.

While I won't be able to fully adjudicate the debate over pragmatic encroachment here, I think we have good reason to resist the claim that such encroachment exists—at least when it comes to

some epistemic standards.¹⁷ For example, whether one has a justified belief (or belief justified to a certain degree) is independent of what level of justification of belief is required for justified *action*. So, in each of the cases above I think it is natural to say that the beliefs based on the evidence are justified to the same degree, no matter the stakes, while at the same time accepting that whether and how one ought to act on those beliefs without acquiring more information is directly affected by the stakes.¹⁸ Compare different standards in the law: for *conviction* of some crimes, belief beyond a reasonable doubt is required; for others, merely a preponderance of the evidence is required. While what counts as ‘beyond a reasonable doubt’ remains the same across contexts, which epistemic standard one needs to meet for some particular action, such as voting ‘guilty’ in a criminal case, can vary. Thus, I believe that the purported cases of encroachment can be resisted, at least when it comes to epistemic standards such as epistemic justification and rationality. It is possible that even if I am wrong about this, there might *still* be a way to separate out two (or more) distinct sorts of responsibility, but I leave this for another time. For now, I want to acknowledge that pragmatic encroachment poses a challenge to the project, but I hope to have said enough to show that it is worthwhile to pursue it.

I now turn to an objection that targets the parallel to the moral when it comes to both aesthetic and epistemic responsibility—and indeed with respect to any domain other than the moral.

VI

The Challenge that Only the Moral Case Meets the Interpersonal Conditions of Responsibility as Accountability. The objection is

¹⁷ It is worth noting that a number of defences of pragmatic or moral encroachment are limited to encroachment as to whether a belief has the status of knowledge (such as Moss 2018), whereas I am more focused on the status of rationality or the degree of justification, for the reason that these standards more naturally lend themselves to responsibility evaluations. I am sceptical about the existence of pragmatic encroachment even in the case of knowledge, for reasons similar to those given in the text, but for present purposes what matters is simply that there are some central epistemic standards that are independent of practical contextual considerations.

¹⁸ In this approach, I am in agreement with Simion (2018), who draws a perspicuous distinction between ‘epistemic norms’ and ‘norms with epistemic content’. The norm of *action* in question might be moral while at the same time having epistemic content. See also Gardiner (2018) for insightful arguments against moral encroachment.

that there is something unique to moral norms that is intimately connected to our practices of blaming for moral transgressions, and it is that moral norms are obligations, and even more importantly, obligations *to* others. In contrast, one might argue, when it comes to pure epistemic norms, either there are no obligations or, if there are, they aren't obligations *to* others. Watson emphasizes the point about the second-personal nature of moral reasons precisely to contrast moral accountability with what he takes to be 'thin' conceptions of moral responsibility or a 'weaker' notion of 'answerability' (Watson 2011, pp. 315–16; see also Watson 2013, p. 473). Reasons are second-personal when they appeal to the legitimate demand of specific others that one not treat them in certain ways. Only a practice that centres on interpersonal norms, or obligations to others, qualifies as a practice of accountability, since accountability is itself an interpersonal notion, involving a relation between the accountable agent and the one to whom she is accountable. If this is right, and if it is also true that when it comes to epistemic norms there are no interpersonal norms of the same kind, then the parallel between moral and epistemic responsibility is limited to responsibility as attributability after all. Or, at the least, there is no sense to be made of epistemic blame except in so far as our beliefs are implicated in moral norms.

To illustrate this reasoning, consider a case that Watson takes to capture a person who is responsible in the attributability sense but not in the accountability sense:

If someone betrays her ideals by choosing a dull but secure occupation in favor of a riskier but potentially more enriching one, or endangers something of deep importance to her life for trivial ends . . . then she has acted badly—cowardly, self-indulgently, at least unwisely. But by these assessments we are not thereby *holding* her responsible, as distinct from holding her to be responsible. To do that, we would have to think that she is accountable to us or to others, whereas in many cases we suppose that such behavior is 'nobody's business'. Unless we think she is responsible to us or to others to live the best life she can—and that is a moral question—we do not think she is accountable here. If her timid or foolish behavior also harms others, and thereby violates requirements of interpersonal relations, that is a different matter. (Watson 1996, p. 231)

Here Watson seems to argue as follows: (i) unless the agent, call her Catrina, owes it to others to act differently—in which case she

would have a moral obligation after all—then she is not accountable to us or others; and (ii) if she is not accountable to us or others, then she is not accountable for these actions. She is responsible in the attributability sense alone. If this reasoning is sound, then it would follow that when it comes to aesthetic responsibility, where we also fail to find obligations to others, there can be no analogue to moral accountability.

Yet I believe that there is reason to doubt both (i) and (ii). Begin with (i). I accept the idea that our practices surrounding moral accountability are essentially interpersonal. The idea of demanding that others live up to their moral obligations necessarily involves both those who demand (even if they are hypothetical demanders) and those of whom such compliance is demanded. But does it follow that the demands themselves must be demands to comply with second-personal obligations, that is, obligations *to* others? It does not seem to me that it does. To see why, consider that I might make a demand of Aida that concerns her treatment of another person, Brett, and hold Aida to account for wronging him. This case illustrates the fact that while accountability is an interpersonal practice, it does not entail that when x holds y responsible in the accountability sense, x holds y to an obligation y has to x . It might be that y has an obligation to z instead. Taking this line of reasoning one step further, there are cases in which x might hold y to an obligation without y 's obligation being to *any* particular person. To see this, consider the fact that people are responsible, and morally blameworthy for, their mistreatment of animals, even ones without very sophisticated cognition, who are not in a position to demand adherence to obligations. It seems clear that people have obligations not to mistreat such animals, even if they do not have obligations *to* the animals themselves in the second-personal sense such that the animals can make demands. Or consider obligations of benevolence. While not universally accepted, it is plausible that we have duties to make others' lives go better when it would not require much in the way of sacrifice, even if others have no right to demand it. Finally, consider the idea that one ought to forgive an offender. It seems to me coherent that in some cases one has an obligation to forgive, without anyone having a right to demand it. In all of these cases, it seems plausible to say that we have moral obligations where no corresponding rights exist. And yet moral accountability does not seem limited to situations only concerning moral obligations that are obligations *to* others. Of

course, one might have a view that the only sorts of moral obligations are of this kind.¹⁹ But the view I just sketched is coherent and at least plausible on its face, and at the same time is consistent with an understanding of the practices of *accountability* as interpersonal. We can hold others to obligations—and more generally to standards—even if they aren't obligations *to us*, or to anyone else. This casts doubt on (i), the claim that accountability, being interpersonal, requires the presence of interpersonal obligations. Further, examining Catrina's case in particular more closely, it is plausible to say that Catrina has duties that are not moral. And she might nevertheless be accountable to her friends, say, who care about her. She might be legitimately on the hook, so to speak, to answer the demands of her friends who blame her for not living up to her ideals.²⁰

Now let us turn to (ii). The fact that Catrina is not accountable to us or to others does not entail that she is not accountable. As long as an agent is *in principle* accountable to others, it seems that she is responsible in the accountability sense. Even in the moral case, I think we can imagine that someone is not accountable to others for some action and yet she can be blameworthy in the accountability sense. In the case of cruelty to animals, for example, it might be that for some reason no one else has the standing to blame or make demands of the offender (perhaps everyone else has been even more cruel, or more radically, everyone else has died). It still seems that the

¹⁹ See Darwall (2006, pp. 28–9) and Scanlon (2008, p. 166) for consideration, if not unequivocal endorsement, of something like this view, and see Zimmerman (2016, pp. 255–6) for an argument against this view.

²⁰ Here I endorse Angela Smith's reaction to the case (see Smith 2015, p. 112). Smith takes it that our moral responsibility practices are not discontinuous with responsibility practices that are non-moral, in a way illustrated by reactions we might have if Catrina were our own friend—we can say 'How could you do this?' and 'I'm so fed up, I can't continue to hang around you', for example—and I am arguing in a similar spirit that we can hold Catrina responsible in this case and blame her in a sense that goes beyond mere attribution of fault. Who has the standing to do this is an interesting question, but I think it is not different from many moral cases in raising it. At the same time, my view diverges from Smith's view in two important ways: first, she is sceptical of the attributability–accountability distinction, whereas I take it that there really is a distinction between two kinds of blameworthiness; second, she wants to capture responsibility with what she calls 'answerability', and in turn requires only that one's rational judgements be reflected in one's actions to count as answerable, and so on her view responsible in the single sense that underlies all our responsibility practices. Nothing like control or opportunity is required on this view, in contrast to the view I endorsed above. See Hieronymi (2008) for a view in a similar spirit to Smith's, according to which moral and *epistemic* responsibility in particular can be seen as accounted for by the same framework that grounds responsibility in a kind of answerability.

offender is blameworthy in the accountability sense, in that demands would be appropriate were someone well positioned to make them.

There is yet a further reason to question the conclusion from (i) and (ii) that accountability requires interpersonal obligations. It is that it makes sense to speak of ‘standing’ to blame in the Catrina case, as it does in uncontroversial cases of moral accountability. Just as many think that it is typically the victim and close friends, say, who have the standing to blame someone like Jon, we might think that Catrina’s close friends can blame her for not living up to her ideals in a way that it would be inappropriate for strangers to do. But questions of standing do not even seem to arise when it comes to responsibility in the attributability sense. Anyone can legitimately make a judgement about moral fault. Thus, there is good reason to think that Catrina is accountable, if not in a moral sense, for her actions, despite the fact that her accountability is not grounded in an interpersonal obligation.

Finally, even if it were the case that there is not enough of the framework of moral accountability to justify the claim that there are genuine analogues in non-moral realms, it is still the case that we go beyond mere attributability in assigning responsibility in other cases. To see this, consider that there is a set of standards that governs the case of Catrina, even though the standards are not given by obligations to others. This is what makes it possible to imagine Catrina’s friend holding her to such standards, and reacting negatively when Catrina falls short. It isn’t that our reaction is merely to offer a negative judgement as we might in the case of a small child’s selfish behaviour. I suggest that Catrina is *deserving* of our negative reaction, or at least she is so on the natural assumption that she *could have* done better. In other words, if we assume that she had a high enough quality of opportunity to meet the standard in question, then we—or at least her friends—could appropriately blame her in a way that goes beyond finding a kind of fault with her.

At this point, an opponent might claim that Catrina’s failure to live up to personal ideals is one kind of case about which an ascription of accountability makes sense, but the specifically aesthetic and epistemic cases are different.

For example, one might try to support this conclusion by arguing that in the blameworthy epistemic cases, say, attitudes such as guilt, resentment, and indignation, which normally accompany or form part of blame in the accountability sense, don’t seem appropriate,

while at least in the Katrina case we can feel something that might be properly described as an analogue of indignation. Is this right? I'm not sure that we are not capable of feeling something analogous to indignation when we think an epistemic agent could have done a better job than she did, given her opportunities. Consider again the case of self-blame in various contexts, such as the inference from polling results or the case in which someone infers that climate change is a 'hoax' based on clearly insufficient evidence. We need to be very careful here, because many cases seem to be 'mixed', in the sense that there are both epistemic and moral considerations. Still, I think that there is some reason to think that epistemic failure, under some circumstances, can be blameworthy in something parallel to the moral accountability sense, and not in virtue of its being a moral failure. It might be that the relevant negative attitudes are not often very similar to resentment and indignation in such cases; but this might simply show that such attitudes are not the defining feature of accountability as has sometimes been thought.²¹

Alternatively, it might be that the general framework applies, but that there are genuine and important specific differences, depending on the domain in question, such as the nature of the relevant attitudes, the conditions and particular persons in particular relationships who are entitled to blame, and more. My aim is simply to show that the same general framework applies in the different domains—a framework that includes at its base the idea of evaluating what one does given one's opportunities in light of applicable standards to support a robust notion of being aptly held responsible.

This raises the question of just how much of the same general framework applies beyond that base. In particular, it is important to address the question of whether the idea of desert applies across domains, or whether it is confined to the moral. Perhaps artists who are aesthetically blameworthy in the accountability sense don't *deserve* for their interests to be undermined in any way if there is not an accompanying moral failure, for example. This is an important question, and I believe that there are at least four possible replies that might be given here. The first is to suggest that desert and accountability can come apart even in the moral case (for example, as [McKenna 2012](#) and [Shoemaker 2015](#) suggest). I have argued against

²¹ Though see [Tollefsen \(2017\)](#) for an interesting elaboration of what she calls 'epistemic indignation'.

this approach elsewhere (Nelkin 2016), but it is worth noting that this approach would have an advantage in being able to provide a maximally unified picture of responsibility if no other good reply can be found. A second approach is to argue that desert is simply a kind of fittingness of negative attitudes, perhaps necessarily painful self-directed ones, and that once we think of desert in this way, it will be more appealing to accept that analogues of moral guilt are indeed fitting across domains. This might be combined with a kind of deflationary understanding of desert as mere appropriateness or fittingness in the way that, say, amusement is fitting in response to a funny joke. Clarke (2013, 2016) and Carlsson (2017) argue that what the blameworthy fundamentally deserve is to feel guilty, and it might be argued that it is appropriate or fitting for the artists or reasoners who do not take their opportunities to do well to experience a negative feeling, even if only a mild one, depending on the circumstances. But it is important that even Clarke and Carlsson argue that there is something special about the kind of fittingness at issue in the moral blameworthiness cases, since other negative feelings in response to states of affairs (such as grief at the loss of a loved one) can also be fitting, and yet seem to be fitting in an importantly different way from desert. For Clarke, what marks the fittingness of guilt in response to blameworthiness specifically is a reason of justice (Clarke 2016), where this is not the case when it comes to grief; for Carlsson what marks the fittingness of guilt is that it provides others with a reason to promote what is deserved, where, again, this is not the case when it comes to grief (Carlsson 2017). Thus, on their view, desert is not simply a generic kind of fittingness which has wide scope and would therefore be uncontroversial to locate in non-moral domains. If they are correct, then, the question of whether desert in the domain of the moral transposes to other domains remains.²²

A third response is to accept the objection and the idea that desert, understood as a special *kind* of fittingness or a special relationship between a person and a valenced response, applies only in the moral case. This is also a live possibility, and if correct, could show that desert and accountability *can* come apart, if only in non-moral cases. To accept this point would constitute a limitation on the unification project, but it would be an interesting one and well worth further

²² In Nelkin (2019) I consider the idea that what is fundamentally deserved in the moral case is guilt.

exploration. But before settling on any of these, it is worth considering a fourth response, which is to simply recognize that desert, along with accountability, is found across domains.

This response leaves many variables to be filled in. For example, the object of desert, or, in other words, what is deserved, may vary greatly across domains. But another dimension that varies across accounts of desert focused on the moral domain is the reason-giving force of desert. Earlier I mentioned that I favour a view that is particularly flexible in this regard. The fact that someone deserves a negative response for doing something morally blameworthy does not by itself provide *any* reason to actually provide that response. Only together with other considerations does the fact that a person deserves something become part of a *pro tanto* reason for their getting what they deserve. In the moral case, it might be that only if it is necessary to distribute some harm does desert become part of a reason to distribute the harm to the deserving first, for example. Thus, even in the moral case, the fact that someone deserves something negative, while having practical import, does not by itself provide a reason. If we could achieve all of the instrumental goods that typically flow from a blameworthy person's getting what she deserves in some other way without cost, then there would be no reason at all to ensure that she gets what she deserves.²³ If this is correct in the moral case, then it is also less difficult to accept the idea that desert too transposes across domains, where it also requires other background conditions to be satisfied before becoming part of a reason to provide a person with what is deserved. It will help to consider an example. By accepting that at least the imagined lazy variant of the artist of *Tomorrow Forever* deserves a negative response, say, we need not thereby commit to even a *pro tanto* reason for us to provide such a response. Thus, the commitments of accepting desert in this domain are less costly than on other accounts of desert in the moral domain. Of course, to accept this account much work would have to be done in articulating the background conditions that *would* make desert part of a reason to provide what is deserved, along with identifying the relevant range of objects of desert. I have only made a start in the moral case, and doing so across domains is a project of its own. Here I hope only to have shown that the defender of a unified

²³ Admittedly, this is quite controversial, and in Nelkin (2019) I defend it against opponents in some detail.

account can appeal to a promising set of resources for addressing the important question of the role of desert across domains, and that, moving in the other direction, whatever the right answer turns out to be when it comes to the degree of similarity across domains could have interesting implications for the proper account of desert in the moral case.²⁴

While we have focused in the last two sections on objections targeting the use of the framework for anything other than moral responsibility, the final objection targets the account as applied in the epistemic case in particular, and poses a special challenge for control views of accountability as suited to a unifying framework.

VII

The Challenge that Only Actions, and not Beliefs, Are Within Our Control. One objection that has often been made to the entire idea of epistemic responsibility is that beliefs are not actions, and not within our control. Thus, they are not candidates for responsibility, blameworthiness, or praiseworthiness. A number of responses have been given to this objection, including those that deny that control is necessary for responsibility in general and those that deny that beliefs are not within our control.²⁵ Given that I have embraced a conception of moral responsibility rooted in opportunity and

²⁴ Writing in the introduction to the second edition of a classic collection of papers on free will about Pettit and Smith's (1996) attempt to offer a unified picture of freedom of action and freedom of thought, Watson notes that the project is important because '[i]t might reorient our views of freedom of the will and moral responsibility. For seeing free will and moral responsibility as a more generic notion removes the emphasis from punishment and blame and brings other dimensions of answerability and accountability to the foreground' (Watson 2003, p. 19). (It must be noted that this insight is in some tension with the objection considered in §V raised by Watson's discussion in another text.) My own view is in significant alignment with Watson's in seeing the unifying project as potentially illuminating of moral accountability, and though I am already inclined to move punishment from foreground to background in the moral case, I am not able to see how we could move blame from foreground to background while still focusing on blameworthiness as an essential aspect of responsibility.

²⁵ Categorizing responses is challenging, because theorists use the term 'control' in different (and sometimes multiple) ways, and to this point I have avoided further elaboration on the term, instead focusing on 'opportunity'. But I think that in so far as it makes sense to distinguish 'control' theorists from 'quality of will' theorists, we must be taking control to be the sort of control that one might have over ordinary actions. A very incomplete list of those who would reject the requirement of control so understood include Smith (2015) and Hieronymi (2008) (see note 20 above), as well as McHugh (2017) and Pettit and Smith (1996), who offer a kind of reasons-responsiveness account of responsibility for beliefs, and

control, I offer here a variant of the latter response, one that builds on some tools from the moral realm.

The first thing to note is that even in the moral realm, not all objects of responsibility are voluntary actions or decisions to act or refrain. We often hold people responsible for unwitting omissions (that is, omissions of which one is unaware at the time, and not the result of any decision to omit). And we typically hold people responsible for at least some of the consequences of their actions or decisions. Thus the fact that beliefs are not themselves something we do does not preclude their being objects of responsibility at the outset. A natural way of thinking about beliefs, then, is as consequences of actions or omissions, and we can thus appeal to some tools, including tracing theories and those that make explicit mention of consequences, from discussions of moral responsibility.

On the most well-known such theories, one's moral responsibility for a later action, omission, or consequence traces back to a prior decision or act of agency such that one could foresee at that time the (risk of the) later action, omission, or consequence (see, for example, Fischer and Tognazzini 2009). In previous work, Sam Rickless and I develop a more minimalist account (Nelkin and Rickless 2017). On the view we favour, whether an agent is morally responsible and blameworthy for x at time t_2 depends entirely and solely on whether there was a prior time, t_1 , at which the agent had the *opportunity* to do something that, as she reasonably believed, would significantly raise the likelihood of avoiding x . Note that what is key here is the having of a prior opportunity to avoid the consequence; having such an opportunity does not require any act of decision or agency. Call this the 'Opportunity Tracing view'.

Now, a very natural objection is that the foreseeability condition isn't met when it comes to beliefs. Precisely because they are not something we do, we are passive with respect to them and can't anticipate—typically, anyway—what we will come to believe. We might be able to gather a lot of evidence intentionally, and foresee coming to a belief; but we cannot foresee what belief we will come to. That's the whole point of intentionally investigating something.²⁶

This objection rests on two premisses: (1) the foreseeability

others sometimes known as 'doxastic compatibilists', such as Ryan (2003) and Owens (2000). See Peels (2017, pp. 73–81) for a nice overview of this position.

²⁶ I take it that this sort of objection is what leads Peels (2017), for example, to move to what he calls an 'influence' view of 'responsible' (that is, 'blameless') belief.

condition requires that one (be able to) know or reasonably believe in advance what one will come to believe on condition one acts (or fails to act) in certain ways; (2) in many cases in which we intuitively hold people responsible for their beliefs, one is not able to know or reasonably believe in advance what one will come to believe on condition one acts (or fails to act) in certain ways. Here I would like to accept (2) for the sake of argument, but reject (1).

There is a deep and difficult question in the moral and legal literature, as well, as to what the foreseeability condition comes to.²⁷ There are at least two important questions here. First, must one actually be aware of a risk, or must one merely be *able* to foresee it? Second, what degree of fineness or coarseness of grain does one need to (be able to) foresee a future action or consequence, conditional on what one currently does or does not do? In other words, must one (be able to) foresee a future act of negligence, say, in perfect detail? For example, does one need the ability to see that one's elderly neighbour to the left will slip on the ice outside of one's shop if one does not remove it, or is it sufficient that one foresees that *someone* will (likely) slip? The view I favour is one that is on the first dimension quite demanding: one must have actual awareness in order to have a genuine opportunity to avoid wrongdoing. But on the other dimension, it is quite undemanding. The 'grain' needn't be fine at all. To see this, consider a case of William FitzPatrick's (from [FitzPatrick 2008](#), based on *It's A Wonderful Life*). If an agent, call him 'Potter', had an opportunity to prevent his later unwitting failure to help his fellow townspeople, that opportunity needn't have been one that included his foresight that he was in danger of causing particular individuals' bankruptcies. Rather, his opportunity might involve simple awareness that failure to take steps to check himself now runs a serious risk of disregarding others' interests later.

Articulating the exact degree of fineness or coarseness of grain in the description of what one must foresee (or be able to foresee) to satisfy a reasonable awareness condition on responsibility for consequences remains an unmet challenge in the moral and legal literature, and I do not have a simple answer. But the example of Potter suggests that, whatever the answer, it will not require something extremely fine-grained. And I believe that something similar carries

²⁷ See, for just some examples, [Vargas \(2005\)](#), [Fischer and Tognazzini \(2009\)](#), and [Shabo \(2015\)](#).

over to the case of belief. There is reason to doubt that what is most relevant for culpable belief is the content of the future belief, as opposed to its normative qualities. Is the important question whether someone can foresee acquiring the belief that the answer to the maths problem is 18,695,021 or that she can foresee acquiring a belief about the answer that isn't well supported by a careful calculation? It is plausible that the latter is the more relevant question. To take another example, if, having found myself engaged in self-deceptive thinking when it comes to an important question concerning my children, I have the opportunity to learn from the experience and take steps to avoid doing something similar in the future, but fail to do so, then it seems I am blameworthy when I do something similar in the future, even when the specifics are entirely unpredictable at the earlier time. And indeed I sometimes do blame myself for not seeing something I should have as a parent. I trace that blame to the idea that I had enough awareness of risk at an earlier time and failed to act. This doesn't require any decision or exercise of agency of any kind. But it does require an opportunity not taken. If this reaction is apt, then it makes it more plausible that we have opportunities to avoid acquiring beliefs for which we can then properly be blamed, and the objection to the Opportunity Tracing view as applied to belief here loses much of its force.

It is worth noting that along some dimensions the picture is very demanding, and in others much less so. On the one hand, actual awareness is required for having a genuine opportunity. On the other, what is required is that one *have* an opportunity, not that one have *exercised* agency or control or one's will; and the content of awareness can be quite coarse-grained.

Importantly, things are not obviously different in the moral case. If you are aware that you can do something without great difficulty that will substantially decrease the risk of doing something seriously moral wrong, then, even without knowing any more about what form that wrong would take or when, that seems sufficient to have met an awareness condition. In each realm, then, there is a significant awareness condition. Thus, the Opportunity Tracing view has application in both the realm of belief and the realm of action, preserving a parallel treatment.

Of course, any general account of accountability for attitudes will leave itself open to the possibility that some of our previous practices were either over-inclusive or under-inclusive. That is the potential

cost of offering any general well-motivated set of conditions. The account explored here—the Quality of Opportunity view together with its tracing component, the Opportunity Tracing view—is not unique in this respect. If it were to be the case that some revision in our patterns of holding accountable is required, that is not necessarily a disadvantage to the view. But a full assessment of the view will ultimately include a comparative evaluation, and here I have simply tried to show that a particular kind of control view that appeals to opportunities has significant resources to offer a unifying picture of responsibility across domains, including the moral and the epistemic.²⁸

VIII

Conclusion. The framework I have sketched recognizes continuity, but not perfect symmetry, among moral and aesthetic and epistemic responsibility and desert. The aspects that are continuous, though, are ones that implicate free will and control—understood as having opportunities to meet or exceed standards to which others could in principle hold us. This gives us one way of capturing an important refrain in work on free will, namely that moral responsibility is one kind among many. And more importantly, I believe it accommodates our practices that incorporate different types of blame and praise. Finally, whether the framework sketched above is correct or not, I am hopeful that continuing to explore the prospects of a unified view promises a better understanding of every particular type of responsibility, including the moral itself.²⁹

²⁸ In work in progress, I compare the virtues of the view set out above to three sets of alternatives. Two such sets share with the Quality of Opportunity view set out above the promise of a unifying account. One set that shares that promise is composed of quality of will views (for example Smith's and Hieronymi's), while the other is a set of control views that are less demanding than the Quality of Opportunity view in not requiring awareness, even of a coarse-grained sort, for opportunity (see, for example, [Clarke 2014](#)). A third set of views might be called 'Influence' views, in that what grounds responsibility for beliefs is some earlier influence on (or failure to influence) the process by which one comes to those beliefs. There are distinct versions of this, including [McCormick \(2011\)](#), [Peels \(2017\)](#), and [Rettler \(2018\)](#), which differ on what the nature of the relevant influence is. Some proponents of influence views allow for a unifying account, but some explicitly suggest that an influence condition is a less robust control condition than is in place for action.

²⁹ Ancestors or parts of this paper were delivered to audiences at the University of St Andrews, Northwestern University, the University of Oslo, GRIN at the Université de Montréal and Fillosophie UQAM, the Florida State Graduate Conference, Ludwig-

Department of Philosophy
University of California, San Diego
 9500 Gilman Drive #0119
 La Jolla, CA 92093-0119
 USA

REFERENCES

- Alston, William P. 1988: 'The Deontological Conception of Epistemic Justification'. *Philosophical Perspectives*, 2: *Epistemology*, pp. 257–99. Reprinted in his *Epistemic Justification: Essays in the Theory of Knowledge*, pp. 115–52. Ithaca, NY: Cornell University Press, 1989.
- Archer, Alfred, and Ware Lauren 2017: 'Aesthetic Supererogation'. *Estekia: The Central European Journal of Aesthetics*, 54(1), pp. 102–16.
- Bradford, Gwen 2015: *Achievement*. Oxford: Oxford University Press.
- Brink, David O., and Dana Kay Nelkin 2013: 'Fairness and the Architecture of Responsibility'. In David Shoemaker (ed.), *Oxford Studies in Agency and Responsibility*, vol. 1, pp. 284–313. Oxford: Oxford University Press.
- Brown, Jessica 2020: 'What is Epistemic Blame?' *Noûs*, 54(2), pp. 389–407.
- Canady, John 1964: 'World's Fair Pavilion Selects Theme Painting'. *New York Times*, 21 February 1964.
- Carlsson, Andreas Brekke 2017: 'Blameworthiness as Deserved Guilt'. *Journal of Ethics*, 21(1), pp. 89–115.
- Clarke, Randolph 2013: 'Some Theses on Desert'. *Philosophical Explorations*, 16(2), pp. 153–64.
- 2014: *Omissions: Agency, Metaphysics, and Responsibility*. New York: Oxford University Press.
- 2016: 'Moral Responsibility, Guilt, and Retributivism'. *Journal of Ethics*, 20(1–3), pp. 121–37.
- Clifford, W. K. [1876] 1999: 'The Ethics of Belief'. In his *The Ethics of Belief and Other Essays*, pp. 70–96. Amherst, MA: Prometheus Books.

Maximilians-Universität München, the University of Texas, and the University of Glasgow, and I am very grateful for feedback on those occasions, as well as from participants in a seminar on the topic I led at UC San Diego. Many thanks to William Albuquerque, Helen Beebe, Monika Betzler, Gunnar Björnsson, David Brink, Jessica Brown, Andreas Carlsson, Randy Clarke, Kathleen Connelly, Cory Davia, Sophia Efstathiou, Christel Fricke, Clé Grattton, Elizabeth Harman, Alex Kaiserman, Chris Kelp, Guy Longworth, Michael McKenna, Ellie Mason, Ben Matheson, Per Milam, Glen Pettigrove, Adam Piovarchy, Sandy Reiter, Sam Rickless, Jeroen Rijnders, Dave Shoemaker, Mona Simion, Sharon Skare, Joe Slater, Paula Sliwa, Angela Smith, Helen Steward, Matt Talbert, Christine Tappolet, Bas Tönissen, Manuel Vargas, Shawn Wang, and Brian Weatherson. Finally, I am indebted to Holly de las Casas for her input and organization, and to the audience at the meeting of the Aristotelian Society.

- Darwall, Stephen 2006: *The Second-Person Standpoint: Morality, Respect, and Accountability*. Cambridge, MA: Harvard University Press.
- Eaton, Marcia Muelder 2008: 'Aesthetic Obligations'. *Journal of Aesthetics and Art Criticism*, 66(1), pp. 1–9.
- Fischer, John Martin, and Neal A. Tognazzini 2009: 'The Truth about Tracing'. *Noûs*, 43(3), pp. 531–56.
- FitzPatrick, William J. 2008: 'Moral Responsibility and Normative Ignorance: Answering a New Skeptical Challenge'. *Ethics*, 118(4), pp. 589–613.
- Gardiner, Georgi 2018: 'Evidentialism and Moral Encroachment'. In Kevin McCain (ed.), *Believing in Accordance with the Evidence: New Essays on Evidentialism*, pp. 169–95. Cham: Springer.
- Guerrero, Alexander A. 2007: 'Don't Know, Don't Kill: Moral Ignorance, Culpability, and Caution'. *Philosophical Studies*, 136(1), pp. 59–97.
- Haji, Ishtiyaque 2016: *Luck's Mischief: Obligation and Blameworthiness on a Thread*. Oxford: Oxford University Press.
- Hampshire, Stuart 1954: 'Logic and Appreciation'. In William Elton (ed.), *Aesthetics and Language*. Oxford: Basil Blackwell, pp. 161–9.
- Hedberg, Trevor 2014: 'Epistemic Supererogation and Its Implications'. *Synthese*, 191(15), pp. 3621–37.
- Hieronymi, Pamela 2008: 'Responsibility for Believing'. *Synthese*, 161(3), pp. 357–73.
- Kane, Robert 1996: *The Significance of Free Will*. New York: Oxford University Press.
- Levy, Neil 2017: 'The Good, the Bad, and the Blameworthy'. *Journal of Ethics and Social Philosophy*, 1(2), pp. 1–16.
- Lopes, Dominic McIver 2018: *Being for Beauty: Aesthetic Agency and Value*. Oxford: Oxford University Press.
- McCormick, Miriam 2011: 'Taking Control of Belief'. *Philosophical Explorations*, 14(2), pp. 169–83.
- McElwee, Brian 2017: 'Supererogation across Normative Domains'. *Australian Journal of Philosophy*, 95(3), pp. 505–16.
- McHugh, Conor 2017: 'Attitudinal Control'. *Synthese*, 194(8), pp. 2745–62.
- McKenna, Michael 2012: *Conversation and Responsibility*. Oxford: Oxford University Press.
- Matheson, Benjamin, and Per-Erik Milam 2020: 'The Case for Non-Moral Blame'. Paper given at Arizona Workshop in Normative Ethics, University of Arizona.
- Meylan, Ann 2019: 'The Reasons-Responsiveness Account of Doxastic Responsibility and the Basing Relation'. *Erkenntnis*, 84(4), pp. 877–93.
- Moss, Sarah 2018: 'Moral Encroachment'. *Proceedings of the Aristotelian Society*, 118(2), pp. 177–205.

- Nelkin, Dana Kay 2011: *Making Sense of Freedom and Responsibility*. Oxford: Oxford University Press.
- 2015: 'Difficulty and Degrees of Moral Praiseworthiness and Blameworthiness'. *Noûs*, 50, pp. 356–78.
- 2016: 'Accountability and Desert'. *The Journal of Ethics*, 20(1–3), pp. 173–89.
- 2017: 'Moral Responsibility for Unwitting Omissions: A New Tracing Account'. In Dana Kay Nelkin and Samuel C. Rickless (eds.), *The Ethics and Law of Omissions*, pp. 106–29. New York: Oxford University Press.
- 2019: 'Guilt, Grief, and the Good'. *Social Philosophy and Policy*, 36(1), pp. 173–91.
- and Rickless Samuel C. 2014: 'Three Cheers for Double Effect'. *Philosophy and Phenomenological Research*, 89(1), pp. 125–58.
- Owens, David 2000: *Reason Without Freedom: The Problem of Epistemic Normativity*. London: Routledge and Kegan Paul.
- Peels, Rik 2017: *Responsible Belief: A Theory in Ethics and Epistemology*. New York: Oxford University Press.
- Pettit, Phillip, and Michael Smith 1996: 'Freedom in Belief and Desire'. *Journal of Philosophy*, 93(9), pp. 429–49.
- Press, Howard 1969: 'Aesthetic Obligation'. *Journal of Philosophy*, 66(16), pp. 522–30.
- Rettler, Lindsay 2018: 'In Defense of Doxastic Blame'. *Synthese*, 195(5), pp. 2205–26.
- Ronson, Jon 2014: 'The Big-Eyed Children: The Extraordinary Story of an Epic Art Fraud'. *Guardian*, 26 October 2016. Online at <https://www.theguardian.com/artanddesign/2014/oct/26/art-fraud-margaret-walter-keane-tim-burton-biopic>.
- Ryan, Sharon 2003: 'Doxastic Compatibilism and the Ethics of Belief'. *Philosophical Studies*, 114(1/2), pp. 47–79.
- 2015: 'In Defense of Moral Evidentialism'. *Logos and Episteme*, 6(4), pp. 405–27.
- Scanlon, T. M. 2008: *Moral Dimensions: Permissibility, Meaning, and Blame*. Cambridge, MA: Harvard University Press.
- Shabo, Seth 2015: 'More Trouble with Tracing'. *Erkenntnis*, 80(5), pp. 987–1011.
- Shoemaker, David 2015: *Responsibility from the Margins*. Oxford: Oxford University Press.
- Simion, Mona 2018: 'No Epistemic Norm for Action'. *American Philosophical Quarterly*, 55(3), pp. 231–8.
- Smith, Angela M. 2008: 'Control, Responsibility, and Moral Assessment'. *Philosophical Studies*, 138(3), pp. 367–92.
- 2012: 'Attributability, Answerability, and Accountability: In Defense of a Unified Account'. *Ethics*, 122(3), pp. 575–89.

- 2015: 'Responsibility as Answerability'. *Inquiry*, 58(2), pp. 99–126.
- Stack, Liam 2016: 'Here's Lucy! "Scary" Statue Is Replaced by One that Looks Like Her'. *New York Times*, 9 August 2016. Online at http://www.nytimes.com/2016/08/10/arts/design/scary-lucy-statue-is-replaced-on-an-iversary-of-comedians-105th-birthday.html?_r=0.
- Stanley, Jason 2005: *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Stroud, Sarah 2006: 'Epistemic Partiality in Friendship'. *Ethics*, 116(3), pp. 498–524.
- Tollefsen, Deborah Peron 2017: 'Epistemic Reactive Attitudes'. *American Philosophical Quarterly*, 54(4), pp. 353–66.
- Tönissen, Bas MS: 'The Possibility of Aesthetic Accountability'. Unpublished manuscript.
- Vargas, Manuel 2005: 'The Trouble with Tracing'. *Midwest Studies in Philosophy*, 29(1), pp. 269–91.
- Watson, Gary 1996: 'Two Faces of Responsibility'. *Philosophical Topics*, 24(2): Free Will, pp. 227–48. Reprinted in Watson 2004, pp. 260–86.
- 2003: 'Introduction'. In Gary Watson (ed.), *Free Will*, 2nd edn., pp. 1–25. Oxford: Oxford University Press.
- 2004: *Agency and Answerability: Selected Essays*. Oxford: Oxford University Press.
- 2011: 'The Trouble with Psychopaths'. In R. Jay Wallace, Rahul Kumar, and Samuel Freeman (eds.), *Reasons and Recognition: Essays on the Philosophy of T. M. Scanlon*, pp. 307–31. Oxford: Oxford University Press.
- 2013: 'Asymmetry and Rational Ability'. *Philosophy and Phenomenological Research*, 86(2), pp. 467–75.
- Weatherson, Brian 2008: 'Deontology and Descartes's Demon'. *Journal of Philosophy*, 105(9), pp. 540–69.
- Wolf, Susan 1990: *Freedom within Reason*. New York: Oxford University Press.
- 2015: 'Responsibility, Moral and Otherwise'. *Inquiry*, 58(2), pp. 127–42.
- Zimmerman, Michael J. 2016: 'Moral Responsibility and the Moral Community: Is Moral Responsibility Essentially Interpersonal?' *Journal of Ethics*, 20(1–3), pp. 247–63.