

# Thinking Outside the (Traditional) Boxes of Moral Luck\*

DANA KAY NELKIN

## 1. INTRODUCTION: MORAL LUCK AND THE LINE DRAWING CHALLENGE

The very idea of moral luck poses a puzzle. In fact, Bernard Williams, who introduced the phrase “moral luck,” writes that he “expected to suggest an oxymoron” (Williams 1993: 251). As I will understand it here, moral luck occurs when an agent can be correctly treated as an object of moral judgment, despite the fact that a significant aspect of what he is assessed for depends on factors beyond his control (Nagel 1979). Williams (1981) had suggested that the idea of luck—or being outside of our control—is simply inconsistent with the idea of moral assessment. The more specific claim that will be the focus here is that the idea of luck is inconsistent with a particular form of moral assessment, namely, moral blameworthiness and praiseworthiness. And in fact, I will focus on a specific kind of moral blameworthiness and praiseworthiness, namely, that associated with moral responsibility understood as accountability. In this sense of responsibility, we are responsible agents insofar as we are subjects of legitimate moral demands and apt candidates for being held to account when it comes to meeting those demands.<sup>1</sup>

\*I am most grateful to Andrew Khoury and Sam Rickless for their very helpful comments on an earlier draft of this paper.

1. This notion of responsibility was articulated by Watson (2004) in contrast to what he called a different face of responsibility, namely, “attributability”. 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.” In contrast, one is responsible in the accountability sense if it is appropriate to make moral demands of an agent. And our practices of holding one another accountable “involve the imposition of demands on people” (Watson 2004: 273). I elaborate on the key notion of responsibility in Section 4.

DOI: 10.1111/misp.12101

© 2019 Wiley Periodicals, Inc.

Yet despite Williams's expressed expectation, many have taken moral luck to be not only a perfectly coherent concept, but one that is very frequently instantiated. And indeed, it does appear that we morally assess people for what they do (or omit to do) despite the fact that what they are assessed for depends on all sorts of factors that are not in their control. We hold them to account, we blame and praise them, even when it appears that what we blame and praise them for depends on factors not in their control. Nevertheless, a puzzle remains, even if moral luck is not strictly an oxymoron, in that we seem to accept it in our actual practices of moral appraisal while at the same time continuing to recognize a general tension between the idea of luck—understood in contrast to control—and moral appraisal.

More specifically, the puzzle arises because the following principle (the Control Principle) seems intuitively correct, particularly when we limit moral assessment to praiseworthiness and blameworthiness:

(CP) We are morally assessable only to the extent that what we are assessed for depends on factors under our control.

And this principle seems to have a natural corollary, namely, that two people ought not to be morally assessed differently if the only other differences between them are due to factors beyond their control. And yet, our actual patterns of blaming and praising seem to rest on our acceptance of moral luck.

For example, we regularly presume that people are more blameworthy for murder than for attempted murder, even if whether the victim dies or not is outside the control of the would-be murderers. For example, the ambulance carrying the victim might face more or less traffic and this might make the difference in whether the victim lives or dies; a condition typically out of any one person's control. Or two people might be equally negligent in getting their brakes checked, but only one of the pair's brakes fail when a dog darts out into the street and is thereby killed. Again, we regularly judge the unlucky member of the pair as more blameworthy than the other.

And yet our reactions in these cases appear to contradict the Control Principle, and the accompanying sentiment eloquently expressed by Immanuel Kant when he writes of the good will that even if it

should wholly lack the capacity to carry out its purpose—if with its greatest efforts it should yet achieve nothing and only the good will were left (not, of course, as a mere wish but as the summoning of all means insofar as they are in our control)—then, like a jewel, it would still shine by itself, as something that has its full worth in itself. Usefulness or fruitlessness can neither add anything to this worth nor take anything away from it. (Kant 1784 [1998]: 4,394)

Though Kant here speaks of the good will, and not of praiseworthiness specifically, the general idea that moral appraisal is fully captured without reference to effects or lack thereof which might be a matter of “misfortune” is a powerful one.

We already have a puzzle, then, in the tension between this idea and our everyday judgments of moral blameworthiness and praiseworthiness that in turn depend on the results of our actions and omissions. This kind of luck is *resultant luck*—luck in the way one's actions or omissions turn out. If we hold everything except for the results fixed in each of two cases, then each member of the pair has exactly the same intentions, has made the same plans, and so on, while things turn out very differently and so both are subject to resultant luck. If we can correctly offer different moral assessments for each member of the pair, then we have a case of resultant *moral* luck. But that is only one form of luck. In his article directly responding to Williams (1981), Thomas Nagel identifies additional forms of luck.

There is also circumstantial luck, luck in the circumstances in which one finds oneself. Nagel's example here is of Nazi collaborators in 1930s Germany who are condemned for committing morally atrocious acts, even though their very presence in Nazi Germany was due to factors beyond their control (Nagel 1979). Had those very people been transferred by the companies for which they worked to Argentina in 1929, perhaps they would have led apparently exemplary lives. If we correctly morally assess the Nazi collaborators differently from their imaginary counterparts in Argentina, then we have a case of circumstantial moral luck. Or take a different example provided by Michael Zimmerman (2002): George shoots and kills Henry. George's counterpart, Georg, would have done the same to Henry's counterpart, Henrik, except that he sneezed before he could pull the trigger. Thus, he did not even attempt to shoot Henrik. If we judge George and Georg differentially blameworthy, then here, too, we have accepted circumstantial moral luck.

There is also constitutive luck, or luck in who one is, in one's own traits or constitution. These, too, are things that are out of our control. To take another example of Zimmerman's, if Georg would have shot Henrik, but for his own timidity, while George, being similar in most respects except less timid, did shoot Henry, then we have an example of constitutive luck because they differ in their constitution or character traits and dispositions. And again, if we judge them differentially blameworthy, then it appears we have accepted constitutive moral luck.<sup>2</sup>

Nagel concludes by pointing out that if all the factors of our actions are ultimately a matter of luck, then it seems that there isn't anything at all for us to be responsible for. This deepens the puzzle we already faced in a tension between the Control Principle and our everyday attributions of blameworthiness and praiseworthiness. We cannot simply make one clear kind of revision to our practices (by, say, denying resultant moral luck); we would instead have to deny that anyone is blameworthy or praiseworthy at all for anything.

2. Nagel also identifies a fourth type of luck, namely, causal luck, or luck in the causal antecedents of one's actions, which he takes to amount to the classic problem of free will. Some (e.g., Latus 2001) have pointed out that this fourth category appears to be redundant if one thinks the causal antecedents of one's actions are exhausted by one's constitution and circumstances. For present purposes, I will set this category aside for now, and focus on the first three types.

Despite—or perhaps prodded by—Nagel’s expressed despair that the puzzle might have no solution, there have been a variety of responses to it. Some have denied the existence of all three kinds of moral luck; others have accepted all three kinds; and still others have accepted either one or two, while rejecting luck of the other kind or kinds. Each of these responses has tried to draw the line either around or between the three types of moral luck. Here I will proceed by suggesting that we can learn a great deal by rejecting—at least at the outset—the presupposition that a line must be drawn at one of the joints of the three types of luck identified by Nagel.

Before defending and employing this approach, it will be important to set out in a bit more detail extant paradigmatic approaches that draw the line at the joints of the three kinds of moral luck. I take on this task in Section 2, identifying advantages and disadvantages of each. Then, in Section 3, I turn to the suggestion that we shouldn’t assume that drawing the line must be done either around or at the joints between categories of moral luck identified by Nagel and adopted and elaborated by many others since. Even if it turns out that the line should be drawn at a joint, I suggest that being liberated from the presupposition is productive. Whatever our response to the problem of moral luck is, it ought to cohere with the best understanding of the nature of moral responsibility itself; so, in Section 4 I set out what I take to be the most plausible conception of moral responsibility, what I call the Quality of Opportunity conception. In Section 5, I explain how this conception of moral responsibility helps explain why we should draw the line where it should be drawn, namely, between two sets of cases, both of which have been treated as paradigms of circumstantial moral luck. Section 6 concludes.

## **2. CANDIDATE PLACES TO DRAW THE LINE BETWEEN ACCEPTABLE AND NON-ACCEPTABLE MORAL LUCK**

One response is to deny the existence of any kind of luck. Zimmerman (2002, 2006, 2015) develops this strategy in the most detail.<sup>3</sup> His approach begins by distinguishing between scope and degree of responsibility. To see how this works, suppose that George shoots and kills Henry, and Georg, who is otherwise just like George, shoots in the direction of Henrik, but a bird appears in the path of the bullet just before it would have hit Henrik, thereby sparing Henrik’s life. In this case of resultant luck, George is responsible and blameworthy for more things than is Georg (e.g., a death), even though, on Zimmerman’s view, they are blameworthy to the same degree. Further, when it comes to circumstantial luck, it might be that an agent is not blameworthy *for* anything at all. The would-be Nazi who moved to Argentina before the war is not blameworthy for any particular act or omission. But, in Zimmerman’s view, he can still be blameworthy to the same

3. Others have paved part of the path toward this response, without explicitly going all the way. For example, Norvin Richards (1986) offers a similar strategy for denying the existence of circumstantial moral luck.

degree as his counterpart who stayed behind, in virtue of the fact that he *would have* acted in the same way in the same circumstances. Nathan Hanna (2014) calls this the “counterfactual” strategy. On this view, one can explain how we can judge the two would-be Nazis blameworthy to the same degree given that the only relevant difference between them is not in their control, and we can do so on the basis of the truth of the relevant counterfactuals. Thus, the account offers an explanatory story in addition to vindicating the Control Principle. Zimmerman extends the account to constitutive luck, as well. Timid Georg is just as blameworthy as bold George, because had he not been so timid—a trait not in his control—he, too, would have killed another person. Again, Zimmerman denies the existence of constitutive moral luck (with a possible caveat<sup>4</sup>), and offers the truth of the relevant counterfactuals as an explanation of how two people can be equally blameworthy despite the fact that only one is blameworthy *for* something. Thus, despite first appearances, there is no moral luck at all. We are all equally praiseworthy and blameworthy, once everything is taken into account.

A second response to the puzzle is simply to accept the existence of luck of all three kinds.<sup>5</sup> A variety of strategies has been used to support this response. First, some distinguish different notions of “luck” and suggest that the appeal of denying resultant luck, say, is really based on a confusion. In the sense that matters, George does control his killing Henry, even if the difference between George and Georg isn’t something in their control.<sup>6</sup> Second, given that most people tend to accept *some* types of moral luck, a prominent strategy is to shift the burden onto opponents to defend their accepting some and not others. Robert Hartman (2017: 105-07) does this in the greatest detail, offering a version of this strategy that is explicitly analogical. Consider three agents who all form the intention and plan to carry out a murder using a gun. Sneezzy sneezes and so is unable to pull the trigger at the single moment at which he might have done and killed the target; Off-Target pulls the trigger, but the bullet is intercepted by a bird before it can reach its target; and Bulls-Eye pulls the trigger and hits her target. By hypothesis, there *is* circumstantial luck, so, claims Hartman, Sneezzy is less blameworthy than Off-Target, even though she would have pulled the trigger had her allergies not acted up. But given that the parallels between Sneezzy and Off-Target (same intentions, plans, and so on) are similar to the parallels between Off-Target

4. Zimmerman (2006: 575) concedes that *if* some characteristics are essential to the identity of the individuals, then the relevant counterfactuals would indeed be false. In that case, “the role that luck plays in the determination of moral responsibility may not be entirely eliminable.” At the same time, he expresses skepticism that such characteristics as timidity are essential to a person’s identity.

5. See, for example, Moore (1997, 2009) and Hartman (2017). Importantly, Hartman thinks that there are limits to the ways in which results and circumstances can affect one’s blameworthiness or praiseworthiness but accepts instances of each form of luck. For example, Hartman offers the following general principle: “An agent is not praiseworthy or blameworthy for an action or not fully so when features of her circumstances outside of her control either fully or partially damage a capacity that is necessary to exercise moral agency” (2017: 93), and also argues that only foreseeable results can affect one’s blameworthiness. See note 12. Presumably, Moore and others would acknowledge that not all results affect one’s blameworthiness, as well.

6. See Moore (1997).

and Bulls-Eye, we have analogical evidence that Off-Target is less blameworthy than Bulls-Eye. And Hartman goes on to explain that the analogy rests on three distinguishable factors: the agents in both pairs have identical agency in a relevant sense despite being distinguished by luck; the agents have very similar mental states; and agents in each pair bring about morally significant events voluntarily (2017: 106-07).<sup>7</sup> Thus, this strategy takes the line-drawing challenge to bring out a kind of arbitrariness in accepting some, but not other kinds of moral luck.

Finally, a third approach takes on this challenge head-on, drawing a line between different categories of moral luck, accepting one and rejecting two kinds, or rejecting one kind while accepting two kinds. Perhaps the more typical variant is to draw a line between resultant luck on the one hand and circumstantial luck and constitutive luck on the other, rejecting the former and accepting both latter kinds. One might argue that the true locus of responsibility is what one intends (or fails to intend), and that this supports the idea that results can vary without affecting one's blameworthiness or praiseworthiness, whereas one's intentions (or omissions to intend) might depend on many factors outside one's control. For example, Adam Smith (1790 [1976]) argued that

To the intention or affection of the heart, therefore, to the propriety and impropriety, to the beneficence or hurtfulness of the design, all praise or blame, all approbation or disapprobation, of any kind, which can justly be bestowed upon any action, must ultimately belong. (II.iii.intro.3)

And Andrew Khoury (2018) has recently developed this idea in a subtle way, arguing that what we are really responsible *for*, at least when it comes to actions, are willings. On this view, luck can affect *what* we will by providing us with circumstances and constitutive features, but moral assessment of the relevant kind takes as its target the willing itself (and perhaps other mental states constitutive of caring, such as emotions and evaluative judgments). Results cannot be part of what we are blamed or praised for, and therefore, there is no resultant moral luck. This offers an explanatory picture of why a line can be drawn between resultant moral luck and the other kinds.<sup>8</sup>

The other main variant of this approach that draws a line between categories accepts constitutive luck, while rejecting both resultant and circumstantial luck. One motivation for this approach is based on the general appeal of Zimmerman's reasoning, but resists one part of it, namely, that it can be extended to constitutive luck. In other words, one might accept the counterfactual strategy, agreeing that sneezy Georg is just as blameworthy as George in virtue of the fact that he would have shot and killed someone had he not sneezed. But one might reject the idea that the relevant counterfactuals are coherent in the case of constitutive luck. For example, perhaps timid Georg's

7. See also Enoch and Guttel (2010: 376-77).

8. There are several powerful arguments for denying resultant luck, but the picture described in the text provides one particular way of explaining and justifying in a systematic way why the line should be drawn precisely here. See Rivera-López (2016) for a different attempt to do this, and Nelkin (2019) for additional discussion.

timidity, among his other traits, is simply essential to Georg. In that case, it would not make sense to say what Georg would have done *had he been less timid*. It would be like asking how many sides a circle would have if it had sides.<sup>9</sup> This approach stands and falls, then, with the counterfactual strategy, together with a serious qualification to its scope.

Each of these approaches has advantages and disadvantages.

First, consider the view that there is no moral luck of any kind. Proponents of this view acknowledge that it is highly revisionary when it comes to our actual pattern of judgments of blameworthiness and praiseworthiness. But perhaps that is not, by itself, a reason to reject it. Yet, it, too, is counterintuitive in even more direct ways. For example, it is hard to match the intuitive appeal of denying resultant luck when it comes to circumstantial luck and constitutive luck. It is not nearly as intuitive that the German who was transferred to Argentina really is as blameworthy as the one who stayed behind. Adding to Nagel's Nazi case, Hanna (2014: 688) offers the following:

Jenny lives in a stable, idyllic, isolated utopian society. Consequently, she hasn't developed to a sufficient degree all the traits that would dispose her to resist tyranny. Unfortunately, for these reasons and for many of the same reasons the emigrant and others would collaborate, Jenny would collaborate if she were in Nazi-Germany-like conditions. Is she as culpable as a Nazi collaborator? I say no. The circumstances just seem too unlikely, too far removed from her actual circumstances to make her that culpable.<sup>10</sup>

When it comes to this case, too, it seems that if one were to adopt the counterfactual strategy, one would be driven by the principle but go against the intuition about a particular case. Nor is it as intuitive that timid Georg is as blameworthy as bold George. This counterintuitiveness might be a price worth paying if the general approach is overall more appealing than its competitors, and perhaps further reflection could diminish the counterintuitiveness. In the end, our assessment will necessarily be comparative, but this is a serious obstacle to overcome.

Turn next to the view that there is moral luck of all kinds. This view has the advantage of avoiding drawing a line between kinds of moral luck, and it accords at least superficially with some of our practices. But when we are careful to compare otherwise similar protagonists differing only in *results*, many have the reflective intuition that they are *equally* blameworthy. For example, when we hold fixed George and Georg's intentions and plans and all else in the case in which Georg's bullet is caught by a bird, the temptation to say that George is more blameworthy evaporates for many.<sup>11</sup> Thus, there are costs to adopting this view in the form of reflective intuitions about

9. See Nelkin (2019) for further discussion of such counterfactuals.

10. See Hartman (2017: 65) for similar sorts of intuition pumps.

11. Notably, the Model Penal Code has even adopted a rejection of resultant moral luck for the punishment of many crimes, and this seems to depend on recognition that culpability for attempted crimes and successful ones that are otherwise the same is equal in degree.

resultant moral luck, as well as a complete rejection of even qualified versions of the Control Principle.

Consider then the view that draws a line between resultant moral luck on the one hand, and circumstantial luck and constitutive luck on the other, rejecting the former and accepting the latter. This view would seem to capture many of the intuitions that either of the all-or-nothing views fails to capture. But it faces challenges, as well. For example, what exactly counts as results? Is it anything after the beginning of the formation of the intention in cases when there is an intention? And how do we distinguish omissions from their consequences? Further, even this view has some potentially counterintuitive implications. For example, it would seem to accept a difference in blameworthiness between Georg who sneezes and George who does not. And yet, intuitively, for many, including me, it is at least tempting to treat these cases similarly when it comes to *degree*.

Finally, consider the other variation of the drawing of a line, this one between resultant luck and circumstantial luck on one side, which are rejected, and constitutive luck on the other, which is accepted. Like the view that rejects all kinds of moral luck, this view, too, has some potentially counterintuitive consequences when it comes to cases like the would-be Nazis or Jenny.

I have provided only a sketch of some possible challenges, and it is true that none of these challenges is fatal in itself. All of the views have some costs, and we must evaluate each as an entire package. The question that we face in the end is how they rate comparatively. But before we turn to that task, I would like to explore the idea that we allow ourselves freer rein in drawing a line than has previously been given. Each of the suggestions canvassed thus far includes whole categories on either side of the line between accepted and rejected moral luck. In what follows, I suggest that we consider drawing a line not between circumstantial luck and one of the other kinds, but between some kinds of circumstantial luck and other kinds. Even if the best response in the end turns out to be to draw the line between two of the traditional categories, setting to the side the presupposition that we must do so can be instructive.

### **3. COUNTERFACTUALS AS INDICATORS AND THINKING OUTSIDE THE TRADITIONAL BOXES**

We are now in a position to see more clearly an initial motivation for doing so. A starting point is that when it comes to the debate over how to treat circumstantial luck, it is striking that proponents of moral luck acceptance and moral luck denial tend to appeal to quite different cases. Circumstantial moral luck deniers tend to point to cases like sneezy Georg to motivate their denial; those who accept circumstantial moral luck tend to point to cases like the Nazi collaborators or Jenny to motivate their acceptance. Could the answer be to find a principled way to accept both of these intuitions? That is what I will explore in the rest of this paper.

Those who deny the existence of circumstantial moral luck often appeal to intuitions, such as that of Georg who sneezes and thereby fails to carry out



his plant to kill Henrik, eliciting the reaction that he is no less blameworthy than George who does not sneeze and kills Henry. The main explanation offered, as we have seen, is that while Georg is not responsible for even an attempt to kill Henrik (so that the *scope* of his responsibility might be zero or at least much smaller than George's), he is as responsible and blameworthy as George in *degree*, and this *in virtue* of its being true that he would have killed Henrik in the same manner had something quite out of his control been different in the circumstances. In other words, it is in virtue of the truth of a counterfactual that Georg is as blameworthy as George, and on this basis circumstantial luck is denied. If this is right, then the "counterfactual strategy" can be extended to other cases—the would-be Nazis, and timid Georg included—because in their cases, too, such counterfactuals are true. They, too, would have acted similarly to their blameworthy counterparts had a feature of their circumstances or constitution that was not in their control been different (see, e.g., Zimmerman 2002).

Those who accept moral luck often work from the opposite direction. Because it seems that timid Georg and bold George are differentially blameworthy, for example, it cannot be that appeal to the counterfactual strategy is sound at all.<sup>12</sup>

I suggest that each side is on to something correct. The counterfactuals have a role to play, though they are not what it is *in virtue of which* anyone is responsible or blameworthy or praiseworthy. The counterfactuals cannot be that in virtue of which anyone is blameworthy if we do not think Jenny is as blameworthy as her counterpart, for example. But that does not mean that they are irrelevant. Counterfactuals are instead, in some contexts, good *indicators* of what really grounds praise and blameworthiness.<sup>13</sup>

12. See Hanna (2014), who argues in this way and also offers other challenges, as well. See also Hartman (2017: 65), who argues in favor of circumstantial luck in part by appeal to cases of this kind. Interestingly, Hartman appeals to cases like the would-be Nazis and Jenny when accepting circumstantial luck and rejecting the counterfactual strategy. For example, he appeals to the case of mild-mannered Sandra who would have freely organized the deaths of hundreds of thousands of people in a distant possible world; but, having secured the pro-circumstantial luck intuition, he then uses cases like sneezy Georg as the basis of his analogical reasoning that there is not only circumstantial moral luck, but also resultant moral luck. This shift makes a kind of sense. The first kind of case elicits a pro-moral luck intuition, while the latter does not. So, if one can rest one's claim to circumstantial moral luck on the first sort of case, and then assume for the sake of the analogical argument that there is circumstantial moral luck in the second, one neatly appeals to different features of the two kinds of cases for different purposes. The point I wish to make is that these differences are not merely accidental. Having presented the beginning of a challenge for Hartman here, it should also be noted that there is an important way in which his approach and the one I am arguing for are similar. As mentioned in note 6, Hartman offers some detailed ways to draw lines within *each* of the categories of moral luck, accepting some and ruling out others. The differences between us are that he thinks that there is moral luck of *all* kinds and he draws lines in different places.

13. Interestingly, a parallel issue arises in connection to contemporary reasons-responsiveness theories of responsibility, such as that of John Martin Fischer and Mark Ravizza (1998). On their account, for example, an agent is morally responsible only if she is receptive to reasons, where that is understood as follows: in at least some other possible worlds in which there were different reasons for doing otherwise, the agent would recognize them. A question arises whether the truth of the specified counterfactuals is just what it *is* to be receptive to reasons, or whether there is something more fundamental that is being responsible and that makes the counterfactuals true. See Brink (forthcoming) for a discussion of this question.

The fundamental question is *what* the counterfactuals indicate. Answering this question takes us into the debate about the nature of moral responsibility. In the next section, I offer an account of responsibility that puts the notion of opportunity at the center of the stage.

#### 4. RESPONSIBILITY AS QUALITY OF OPPORTUNITY

Return again to the idea of responsibility as accountability, understood fundamentally in terms of the aptness of demands. It seems that in order to be accountable, for it to be appropriate to others to hold one to moral demands, one must have the ability to comply with them. 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.<sup>14</sup> In particular, one must have the opportunity to avoid wrongdoing.

Going a step further, in order to be *blameworthy* in the accountability sense, one must have not only an opportunity to avoid wrongdoing, but also a *high-enough quality* opportunity. These satisfaction conditions on blameworthy action in the accountability sense gain support from the fact that a conception of blameworthiness in terms of the fair opportunity to avoid wrongdoing best captures the wide variety of commonly recognized excuses. 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 one) 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 things for the right reasons.<sup>15</sup>

Importantly, the quality of opportunities is scalar. The worse an opportunity is, because it would be, say, very difficult to act rightly, the more mitigating of blameworthiness if one fails to do so. On the flip side, if one does act well for good reasons in such circumstances, one is more praiseworthy than if one acted well when it is easy to do so.

Now, in grounding responsibility in the nature of one's opportunities to meet relevant demands (and ultimately in the obligations or standards on which these rest), the Quality of Opportunity view is best categorized as a

14. For further elaboration of this sort of view, see Nelkin (2011) and Nelkin (2016). Susan 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 calls "free and responsible" action. While I favor an understanding of these conditions that is compatible with the truth of determinism and have defended that understanding in the works just cited, I will here set aside this issue, and be as ecumenical as possible.

15. For a related discussion of different kinds of excuses, see Brink and Nelkin (2013).

“control” view, in contrast to views that emphasize the nature of evaluative judgment or quality of will expressed in action or attitude. It is essential that one have control in order to be responsible. But what does having control come to? Does it require an act of agency, even if just a mental act of choice?

If so, this would appear to rule out responsibility for certain sorts of omissions, such as forgetting to do what is needed to keep a promise. Many control theorists, wanting to accommodate such responsibility, add to their theories by adopting a so-called “tracing” condition. In other words, even if one doesn’t meet the conditions for control (or, on the account I’ve set out, an opportunity of high- enough quality at the time of the act or omission), one might be responsible for that act or omission in virtue of an earlier moment in which one did meet the relevant conditions and at which the risk of the later act or omission was foreseeable. Thus, perhaps one earlier chose not to set an alert on one’s phone to remind one to keep one’s promise; in that case, one is responsible for the omission in virtue of one’s earlier choice.<sup>16</sup>

At this point, although we have pushed back the question of what is required for an opportunity to be of good enough quality for blameworthiness, the question nevertheless remains. In many cases of so-called “unwitting omissions” such as when one forgets to do something one has promised, it is not at all obvious that one ever exerted one’s agency in one direction or another, or made any sort of choice related to forgetting or remembering. So, control views that require an earlier choice or exercise of agency (even if just a choice not to act) would seem ill-equipped to accommodate many cases of unwitting omissions that intuitively we take to be ones for which people are responsible and blameworthy. In work with Samuel Rickless (Nelkin and Rickless 2017), I have argued that an opportunity of high enough quality, and one in which one has the requisite control, need not be one in which one exercises agency or makes a choice. One simply must have had the opportunity to do so, and this, in our view, requires *awareness* of the risk of not acting.<sup>17</sup> But one might be aware that one could make a choice, and simply not make one. Other things equal, this would be sufficient, on our view, for having had an opportunity of sufficiently high quality such that one could be blameworthy in the future for an omission in virtue of one’s not having taken it. We call this “the Opportunity Tracing view.”

Combining these elements, then, the account of responsibility I propose here is one that takes the quality of opportunity to be central. One can be responsible and blameworthy if one has a high-enough quality of opportunity

16. For some foundational discussions of tracing, see Fischer and Ravizza (1998), Vargas (2005), and Fischer and Tognazzini (2009).

17. More formally: whether an agent is morally responsible for an unwitting omission at time T2 depends entirely and solely on whether there was a prior time, T1, at which the agent had the opportunity to do something that, as she reasonably believed, would significantly raise the likelihood of avoiding later omission (Nelkin and Rickless 2017). Notably, Clarke (2014) provides an example of a Quality of Opportunity view that is even less demanding; not only does it not require an exercise of control, it does not require awareness of risk either. Nevertheless, on his view, one can be blameworthy if one has an opportunity and fails to take it. Thus, on our view, as well as on Clarke’s, no exercise of agency is required for having an opportunity of high-enough quality.

and does not take it at the time of action (or omission), *or* if one earlier had a high-enough quality of opportunity at which time one was aware of the risk of acting (or omitting) badly later. This picture leaves out many details, but my aim here is simply to set out what I take to be an appealing framework of responsibility, blameworthiness, and praiseworthiness, and to see how it can mutually support an appealing response to the puzzle of moral luck.<sup>18</sup>

## 5. THINKING OUTSIDE THE TRADITIONAL BOXES

Now let us return to the debate over circumstantial luck. Recall that being transferred to Argentina from Germany and sneezing before one can form the intention to shoot at a target are both typically categorized as examples of circumstantial luck. And we can form counterfactuals about both that bring this out: had the German not been transferred, he would have cooperated with the Nazis; had Georg not sneezed, he would have shot Henrik. Each is claimed by deniers of circumstantial moral luck to be as blameworthy as their counterparts (e.g., Zimmerman), while each is claimed by accepters of circumstantial moral luck to be less blameworthy (e.g., Hartman). But if we allow ourselves to treat these cases differently, and can find a principled way of doing so, we can vindicate differential intuitions about the two pairs of cases. And I believe that the Quality of Opportunity conception of responsibility provides just the way to do this.

As we have seen, the counterfactuals can have a role to play, even though they are not what it is *in virtue of which* anyone is responsible or blameworthy or praiseworthy. They are instead, in some contexts, good *indicators* of what really grounds praiseworthiness and blameworthiness, namely, what people's opportunities really are and what they do with them. If we know that Georg would have killed Henrik but for the sneeze, then we know (or have excellent evidence) that in the circumstances he had formed the firm proximal intention to kill him when he had the opportunity to do so. Arguably the blameworthiness in both Georg's and George's cases derives ultimately from this moment. Both are equally responsible because they had a morally similar opportunity and took it in the same way. It is crucial to note at the same time that the use of counterfactuals as indicators has its limits. The truth of counterfactuals in the case of the would-be Nazi or Jenny

18. Andrew Khoury (in correspondence) raises the question of the relationship between this account and so-called "Quality of Will" views of accountability and blameworthiness, and whether one can be seen as a version of the other. I think that the answer is not straightforward. On the one hand, with the emphasis on control and opportunities to act (even if not taken), they seem to be opposed to views that take acting or omitting with a certain quality of will to be *sufficient* for blameworthiness (regardless of opportunities to avoid so willing or acting or omitting). See, for example, Talbert (2008). At the same time, I think that there is a way of understanding "quality of will" so that it is a *function* of what one does or does not do with one's opportunities. The fact that the Quality of Opportunity view does not require an actual exercise of control might make this more attractive than a "control" view that requires an exercise of agency for blameworthiness insofar as it allows for agents to express quality of will even when there is no exercise of agency.

or timid Georg does not point to opportunities that are similar in moral quality to their counterparts. Once we see that the legitimate role of counterfactuals is that of an indicator of that which is fundamentally explanatory, we can see that only some counterfactuals play such a role. Given this, we have a principled way of treating the two cases differently.

Putting the pieces together, then, we can see that the Quality of Opportunity view offers grounds for treating some cases that have been categorized as circumstantial luck as genuine moral luck, while not so treating others.

At this point, one might think that the real line ought to be drawn between cases in which there are intentions or plans and ones in which there are not. In this case, the counterfactuals point us toward the formation of an intention or a plan or a willing, and this can support the weight of the attribution of equal blameworthiness in each member of the pair.<sup>19</sup> This might recall the spirit of Adam Smith's suggestion, defended by Khoury, that draws the line between resultant luck and the other sorts of luck. Further, this approach might be taken in a direction that simply categorizes some cases initially categorized as circumstantial luck as in fact resultant luck. We might think that what really matters is the formation of an intention or a plan or a willing, and that all that follows (including the actual attempt) is a result for relevant purposes.

This is an elegant solution, but I believe that the picture needs some modification. Consider omissions, for example. Imagine two agents, both of whom have promised to pick up the neighbor children from school with their own. One forgets and leaves the neighbor children behind. The other remembers, but only because, as he is scanning email in the parking lot, he happens to see a message from someone whose name is similar to that of his neighbors' and it jogs his memory. The fact that the second driver is lucky in a way, indicated by the counterfactual that he would have forgotten but for the chance reminder, seems important. He, too, had forgotten, and this seems an important fact. If we think that there was an opportunity of high-enough quality at some point along the way that the first driver failed to take, then it seems that the second driver also had such an opportunity and failed in the same way. So, this counterfactual helps indicate the existence of such a similar opportunity not taken. But in the first case, there was no intention or plan to act wrongly; the very idea of an unwitting omission like this one entails that there was no such thing. And yet, intuitively, this too seems like a case where it would be inapt to blame the first driver more than the second.<sup>20</sup>

19. See Enoch and Marmor (2007) for consideration but ultimate rejection of this hypothesis, along with rejection of both resultant and circumstantial luck more generally.

20. One might point out that the second driver had an additional opportunity not available to the first driver. And that is true. But in this case, the second driver did no more than one could reasonably expect at low cost, so it is plausible that it is not praiseworthy. If we change the case, though, so that what the person was expected to do was quite demanding or required great sacrifice, then it seems that the second agent, upon his memory being jogged, really did have an opportunity that was saliently different from the first, but also that the second is intuitively more praiseworthy if he followed through.

Or imagine a case in which no act of agency is manifested at any point, not even in the giving of a promise. A driver sees someone whose car has broken down in a pre-cell phone era. She is aware that she could call for help at the next stop, does not make a decision either way, but continues to think about the problem of moral luck, or perhaps thinks about nothing at all. She passes the next stop and never again thinks about the person on the side of the road. No willing need be involved at all. And yet, the driver is blameworthy precisely for not having taken steps to help, and this seems true regardless of what happens to the person on the side of the road, or whether her thoughts are joggled by a fluke message on the radio later.<sup>21</sup>

A more comprehensive unifying explanation, then, is that the counterfactuals are important indicators of opportunities taken or missed, and it is this that is key in determining degree of blameworthiness, rather than any particular act of will.

Thus, there is in fact moral luck in generating what opportunities one faces. Yet, once one has opportunities—and opportunities of a high-enough quality—then one is properly assessed only on the basis of what one does, or fails to do, with them. One’s responsibility and degree of blameworthiness or praiseworthiness are a function of the quality of one’s opportunity and what one does with it. And importantly, where there is luck, not all of it is bad. It might be that some people are lucky enough to be in a position to perform praiseworthy acts while others who would have done so had they been in that position are not.<sup>22</sup>

After having filled out the picture in this way, it might seem that it can be fit into one of the traditional approaches after all: it draws a line between circumstantial and resultant luck, accepting the first and denying the second. There is a way of understanding “resultant” luck in which I think that this is true after all, namely, that results are the way the world unfolds once one’s opportunity and one’s response (or nonresponse) is fixed. At the same time, giving ourselves the freedom to distinguish cases that have traditionally been grouped together as circumstantial on the basis of sharing a certain sort of counterfactual feature allows us to recategorize the cases. Cases that appeared to be candidate cases of circumstantial luck are instead candidate cases of

21. The difference between this view and the version of Khoury’s view on which it is willings that we are ultimately responsible for is thus subtle, but important. Khoury points out that one attractive feature of the picture on which it is willings that are fundamental is its association of responsibility with control. On the view I have presented, the focus on willings goes too far, however, in associating responsibility with the *exercise* of control. On the Quality of Opportunity view, what is crucial is that there be the opportunity for such an exercise. But, as in the case of the driver who thinks about nothing at all instead of taking the opportunity to make a mental note to call for help, there might be no such exercise, even though there is an opportunity. Given that blameworthiness even in that case seems apt, a picture according to which what is fundamental is what one does or fails to do with one’s opportunities seems better able to capture the cases. And to the extent that we value and assess on the basis of control, it seems enough to *have* control in the sense of having an opportunity, whether or not one exercises it.

22. I explore the special case in which one person is lucky enough to have the features that make her an apt candidate for responsibility and praiseworthiness, while some are not, in Nelkin (2011: 39).

resultant luck on this expansive understanding of resultant luck. Thus, even if, in the end, the best approach is to draw a line between resultant luck and the other kinds, we will have made progress in two ways by having freed ourselves from the presupposition that there are only a very limited number of places to draw the line. The first way is that it allows us to think about cases that have been classified together quite differently, and the second is that it allows us to think outside the box about the nature of the starting point from which we then measure what comes after as results.<sup>23</sup>

I believe that this response to the problem of moral luck and the Quality of Opportunity conception of moral responsibility are mutually supporting and allow us to preserve more intuitive reactions to both cases and principles than others. At the same time, it is worth noting that the general approach of freeing ourselves from the presupposition that the line must be drawn around or between the boxes can also be adopted by competing conceptions of moral responsibility.<sup>24</sup>

23. I note that there will be more complicated cases than those discussed in the text, such as cases in which agents form plans and abandon them before they can be carried out. While I believe that such cases can be handled in intuitive ways by the Quality of Opportunity account developed here, they would need a more extensive treatment than I can provide here.

24. On a competing framework, the relevant counterfactuals, when good indicators, are not indicators of opportunities taken or missed, but rather indicators of a person's character, or dispositions, or evaluative commitments. These are all very different items, but they all seem to fall under an umbrella of "Quality of Will." Consider, to take just one case that carves things in a clearly different way, Richards's (1986) classic argument against circumstantial luck. He takes it that the locus of our responsibility judgments is character, where character must be enacted in some way, if only in one's mental life. As he writes, we do care if our friend would have left us in the lurch when we needed her, even if we didn't in fact find ourselves with the need. And the counterfactual that she would have acted in that way makes us blame her and judge her as blameworthy as if she had acted in such a way when we needed her, even if, as it happens, she did nothing at all. On this view, it is their character, and the accompanying orientation of their will toward us, that makes the counterfactuals true and that is doing the real work. Though Richards does not consider a different variant of the Nazi case here, and perhaps he does not think it is coherent, we might imagine one in which the relevant counterfactuals are true, and yet the agent does *not* have the relevant character or orientation of will. Perhaps he has the proto-dispositions, such that *were* he to be in Germany, these other dispositions would be activated in some sense. But as it is, he isn't properly described as having such a strong desire to please authorities or be unaverse to suffering. In that case, a proponent of this "character" view of what grounds desert and blameworthiness should say that the truth of the relevant counterfactuals *does not* in fact make the transplanted German blameworthy to the same degree. On this view, some of the relevant kinds of counterfactuals will point to a shared feature that undergirds equal blameworthiness, while some will not. We can see this as a view that splits the difference in a principled way between different kinds of circumstantial luck, or we can see this as drawing the line between circumstantial luck and constitutive luck after all. But even if the latter were true, dispensing with the idea that cases originally all thought to be fit into the circumstantial luck category have to be treated similarly, and giving up the idea that all counterfactuals that advert to circumstances are to be treated on a par, allows for a different categorization, driven fundamentally by a conception of the nature of responsibility. This kind of view has not been at the center of the debate about moral responsibility in recent years, in part due to the influence of situationist moral psychology that calls into question the existence of global, stable character traits, but it has been influential in the literature in legal theory. See Moore (2009: chapter 13) for an excellent overview and evaluation. Other versions of Quality of Will views, such as the view of Angela Smith (2015), that takes evaluative judgments as central, will no doubt have very different implications.

## 6. CONCLUSION

Interestingly, many who write about moral luck tend to focus on cases of blameworthiness. But it is notable that in the famous passage that Nagel cites before going on to offer a number of cases in which people act badly, Kant is focused on the *good* will, one whose best efforts might come to nothing, but which nevertheless shines like a jewel. Though I, too, have here focused largely on cases of blameworthy action, I believe the picture I've presented dovetails with Kant's vivid imagery and ideas. If the ultimate ground of our responsibility is a combination of our opportunities and what we do—or don't do—with them, then if we fail to will the good—by willing badly, willing indifferently, or not willing at all—when we have an opportunity of high-enough quality to do so, then we are blameworthy. But if we seize our opportunity, and we have done *our* best in exerting great effort toward the good, then we may be praiseworthy, even if the outcome does not match our will.

## REFERENCES

- Brink, David O. Forthcoming. *Fair Opportunity, Responsibility, and Excuse*. Oxford: Oxford University Press.
- Brink, David O., and Dana Kay Nelkin. 2013. "Fairness and the Architecture of Responsibility." *Oxford Studies in Agency and Responsibility* 1: 284–313.
- Clarke, Randolph. 2014. *Omissions: Agency, Metaphysics, and Responsibility*. New York: Oxford University Press.
- Enoch, David., and Ehud Guttel. 2010. "Cognitive Biases and Moral Luck." *Journal of Moral Philosophy* 7: 372–86.
- Enoch, David, and Andrei Marmor. 2007. "The Case Against Moral Luck." *Law and Philosophy* 26: 405–36.
- Fischer, John Martin, and Mark Ravizza. 1998. *Responsibility and Control: A Theory of Moral Responsibility*, Cambridge: Cambridge University Press.
- Fischer, John Martin, and Neil A. Tognazzini. 2009. "The Truth about Tracing." *Noûs* 43: 531–56.
- Hanna, Nathan. 2014. "Moral Luck Defended." *Noûs* 48: 683–98.
- Hartman, Robert. 2017. *In Defense of Moral Luck: Why Luck Often Affects Praiseworthiness and Blameworthiness*. Abingdon, UK: Routledge.
- Kant, Immanuel. 1784 [1998]. *Groundwork of the Metaphysics of Morals*. Mary Gregor (ed. and trans.), Cambridge: Cambridge University Press.
- Khoury, Andrew. 2018. "The Objects of Moral Responsibility." *Philosophical Studies* 175: 1357–81.
- Latus, Andrew. 2001. "Moral Luck." *The Internet Encyclopedia of Philosophy*. J. Fieser and B. Dowden (eds.). Retrieved February 1, 2019, from <https://www.iep.utm.edu/moralluc/>.
- Moore, Michael. 1997. *Placing Blame: A Theory of the Criminal Law*. Oxford: Clarendon Press.
- . 2009. *Causation and Responsibility: An Essay in Law, Morals, and Metaphysics*. Oxford: Oxford University Press.
- Nagel, Thomas. 1979. *Mortal Questions*. New York: Cambridge University Press.
- Nelkin, Dana Kay. 2011. *Making Sense of Freedom and Responsibility*. Oxford: Oxford University Press.
- . 2016. "Difficulty and Degrees of Moral Praiseworthiness and Blameworthiness." *Noûs* 50: 356–78.



- . 2019. “Moral Luck.” The Stanford Encyclopedia of Philosophy (Summer 2019 Edition), Edward N. Zalta (ed.). Retrieved May 1, 2019 from <https://plato.stanford.edu/archives/sum2019/entries/moral-luck/>.
- Nelkin, Dana Kay, and Samuel C. Rickless. 2017. “Moral Responsibility for Unwitting Omissions: A New Tracing Account.” In *The Ethics and Law of Omissions*, Dana Kay Nelkin and Samuel C. Rickless (eds.), 106–29. New York: Oxford University Press.
- Richards, Norvin. 1986. “Luck and Desert.” *Mind* 65: 198–209.
- Rivera-López, Eduardo. 2016. “How to Reject Resultant Luck Alone.” *Journal of Value Inquiry* 50: 415–23.
- Smith, Adam. 1790 [1976]. *The Theory of Moral Sentiments*. D. D. Raphael and A. L. Macfie (eds.). Oxford: Clarendon Press.
- Smith, Angela M. 2015. “Responsibility as Answerability.” *Inquiry* 58: 99–126.
- Talbert, Matthew. 2008. “Blame and Responsiveness to Moral Reasons: Are Psychopaths Blameworthy?” *Pacific Philosophical Quarterly* 89: 516–35.
- Vargas, Manuel. 2005. “The Trouble with Tracing.” *Midwest Studies in Philosophy* 29: 269–91.
- Watson, Gary. 2004. “Two Faces of Responsibility.” In *Agency and Answerability: Selected Essays*, 260–88. Oxford: Oxford University Press.
- Williams, Bernard. 1981. *Moral Luck*. Cambridge: Cambridge University Press.
- . 1993. “Postscript.” In *Moral Luck*, Daniel Statman (ed.), 251–58. Albany: State University of New York Press.
- Wolf, Susan. 1990. *Freedom Within Reason*. Oxford: Oxford University Press.
- Zimmerman, Michael. 2002. “Taking Luck Seriously.” *The Journal of Philosophy* 99: 553–76.
- . 2006. “Moral Luck: A Partial Map.” *Canadian Journal of Philosophy* 36: 585–608.
- . 2015. “Moral Luck Reexamined.” *Oxford Studies in Agency and Responsibility* 3: 136–58.