

Moral Principles: A Challenge for Deniers of Moral Luck

Anna Nyman (Uppsala University)

[This is a penultimate draft. Please cite the final version forthcoming in *Ergo*.]

Abstract. On a common characterization, moral luck occurs when factors beyond agents' control affect their moral responsibility. The existence of moral luck is widely contested, however. In this paper, I present a new challenge for deniers of moral luck. It seems that some factors beyond agents' control – such as moral principles about blame- and praiseworthiness – clearly affect moral responsibility. Thus, moral luck deniers face a dialectical burden that has so far gone unnoticed. They must either point to a relevant difference between factors like moral principles and the kind of factors that according to them do not affect moral responsibility or show how they can avoid having to point to such a difference. I argue that no obvious way to meet the challenge presents itself and that it thus amounts to a serious worry for deniers of moral luck.

1. The Problem of Moral Luck

On a common characterization, moral luck occurs when an agent's moral responsibility – her moral blame- and praiseworthiness – is affected by factors that are beyond her control. A familiar understanding of the problem of moral luck has it that it consists in a contradiction to which common-sense ideas about moral responsibility seem to commit us: that moral luck exists and that it does not.¹

The usual setup of the problem is the following. Common-sense moral responsibility judgments in particular cases indicate that we are committed to the existence of moral luck. Common sense has it, for instance, that assassins who manage to kill their targets are more blameworthy than assassins who fail to do so, that assassins who get as far as committing a murder attempt are more blameworthy than assassins who do not, and that someone who grows up to become the kind of person that takes an offer to murder in exchange for money is more blameworthy than someone who grows up to become the kind of person that refuses the offer in the first place. If that is so, moral luck exists because these agents' moral responsibility is then affected by factors beyond their control. Consider the following scenarios to see this:

¹ In this form, the problem originates with Nagel (1979/2012). For more background, see e.g. Nelkin (2021).

Results. Killer and Birdy are both hired assassins and take a shot at their respective target. Only Killer hits and kills her target, however, since a passing bird intercepts Birdy's bullet.

Circumstances. Birdy and Sneezzy are both hired assassins and they plan to take a shot at their respective target. Only Birdy gets an opportunity to do so, however, since some pollen in the air makes Sneezzy sneeze at a crucial moment.

Constitution. Saint and Sneezzy are both offered to kill a target in exchange for money. Only Sneezzy takes up the offer, however, because only she had an upbringing and a genetic makeup that resulted in a greedy and violent disposition.

If Killer is more blameworthy than Birdy, whether a bird intercepts one's bullet affects moral responsibility. If Birdy is more blameworthy than Sneezzy, whether there is pollen in the air affects moral responsibility. And if Sneezzy is more blameworthy than Saint, whether one has a certain genetic makeup and has had a certain upbringing affect moral responsibility. Thus, since all of these factors are beyond the agents' control, we are committed to the existence of moral luck if we judge them blameworthy to different degrees.

The particular cases where our judgments indicate a commitment to moral luck are usually divided into different categories, where the sorting rationale is the type of factor beyond agents' control that affects moral responsibility. If it is a resultant luck factor – one that affects how their action turns out, like in *Results* – we have a case of *resultant* moral luck. If it is a circumstantial luck factor – one that affects what opportunities and choices agents face, like in *Circumstances* – we have a case of *circumstantial* moral luck. And finally, if it is a constitutive luck factor – one that affects what dispositions and character traits an agent has, like in *Constitution* – we have a case of *constitutive* moral luck.²

On the one hand, then, we have what seems like a commitment to moral luck. On the other hand, common sense also seems committed to the so-called *control principle*, that is, to the idea that agents are responsible for something only to the extent that it depends on factors within their control.³ Many find this principle highly intuitive; some call it “a bedrock assumption in our ordinary moral thinking” (Hartman 2017: 4). But here is the crux: if the control principle is true, it

² A fourth category, causal moral luck, is also acknowledged in the literature. Causal luck has, according to Nagel (1979/2012: 28), to do with “how one is determined by antecedent circumstances”. If such a factor affects an agent's moral responsibility, there is causal moral luck. This kind of moral luck is often bracketed in the debate, however, and I follow suit.

³ This particular phrasing of the principle is from Nelkin (2021).

is held, then there can be no moral luck – in that case, agents like those in the scenarios above must be equally blameworthy since the difference between them is due to factors beyond their control.⁴

It is thus thought that we have to give up either the particular moral responsibility judgments or the control principle to solve the problem of moral luck. The most popular strategy is to abandon the particular moral responsibility judgments (e.g., Enoch and Marmor 2007; Herdova and Kearns 2015; Jensen 1984; Khoury 2018; Levy 2011; Lockhart and Lockhart 2018; Nelkin 2019; Peels 2015; Pritchard 2005; Rescher 1993; Richards 1986; Rivera-López 2016; Rosebury 1995; Sverdlik 1988; Thomson 1989; Zimmerman 2002; 2015). Some argue that we should do so entirely, by denying all moral luck, while others argue that we should do so partly, by denying some kinds of moral luck and accepting other kinds. In this paper, I put forth and discuss a challenge for both sorts of moral luck deniers.

2. The Challenge

As noted, a common characterization of moral luck has it that moral luck occurs when an agent's moral responsibility is affected by factors beyond her control. I believe, however, that consistently denying moral luck is harder than previously recognized, because there seems to be little room for denying that certain factors are both beyond agents' control and affect moral responsibility.

Consider, for instance, the truth of principles about moral responsibility and the deontic status of actions. It is beyond an agent's control that a correct moral principle condemns a certain action of hers, and yet the principle's condemning it surely affects her moral responsibility in that it is because the principle is true that she will, given that she fulfils conditions for moral responsibility, be blameworthy for the action rather than praiseworthy. Likewise, it is beyond an agent's control that a correct moral principle says that agents who act as she does are blameworthy for what they do rather than praiseworthy. Yet the principle's saying so plainly affects her moral responsibility in that it is because the principle is true that she is blameworthy rather than praiseworthy. Thus, since moral principles are both beyond agents' control and affect moral responsibility, it would seem that quite a bit of moral luck exists.⁵

The idea that it matters for the controversy over moral luck that factors such as moral principles affect moral responsibility may at first seem incredible, and deniers of moral luck might be tempted to put it aside as trivial that they do. "What we deny", they might say, "is not *that*, but rather that

⁴ Hanna (2014) argues, however, that the control principle is compatible with the existence of moral luck. See Zimmerman (2015) for a reply.

⁵ One may wonder whether moral principles fit within any of the four categories of luck factors outlined in the previous section. Moral principles may classify as circumstantial luck factors, thus giving rise to circumstantial moral luck. But whether they do is a question we can put aside since there is no particular reason to think that these categories are exhaustive.

standard luck factors, the kind of luck factors present in standard moral luck scenarios, affect moral responsibility – that is, factors such as whether a bird intercepts the bullet one has fired at a target, whether the presence of pollen in the air makes one sneeze, and whether one’s upbringing and genetic makeup result in a greedy and violent disposition.”

It amounts to a serious worry for deniers of standard moral luck, however, that certain non-standard factors beyond agents’ control affect moral responsibility. This is because it reveals that these deniers face a dialectical burden that has so far gone unnoticed. They must either point to a relevant difference between non-standard luck factors like moral principles and the kinds of standard luck factors that according to them do not affect responsibility (call this *the relevant difference task*), or find a way to avoid having to point to such a difference. Only when standard moral luck deniers have done so are they justified in maintaining their position.⁶

In the rest of this paper, I explore the prospects of dealing with the relevant difference task. For simplicity, I focus on one example – moral principles – but it is worth keeping in mind throughout that to deal with the relevant difference task in its entirety, deniers of standard moral luck must attend to each factor that seemingly trivially gives rise to moral luck. Other factors include, for instance, that one is born (cf. Hanna 2014: 695, fn. 2), and that one has the mental capacities to be an agent at all.

I will first consider whether deniers of standard moral luck can avoid (this version of) the relevant difference task by denying that moral principles give rise to moral luck. They could do so either by claiming that the standard characterization of moral luck is faulty (a response I consider in section 3) or by insisting that moral principles do not, ultimately, affect moral responsibility (a response I consider in section 4).

After establishing that no obvious way to avoid the relevant difference task presents itself, I move on to look at two ways in which it might potentially be solved. A seemingly obvious reply is to say that moral principles are necessary, and that this fact will surely allow one to somehow solve the challenge. I consider (in section 5) the problems with appealing, directly or indirectly, to necessity as the way in which moral principles relevantly differ from standard luck factors. I then look at a second potentially relevant difference, namely, that while standard luck factors are matters of luck, moral principles are not (section 6).⁷

⁶ It is worth noting that the relevant difference task can be seen as an aspect of what Nelkin (2021) calls the line-drawing challenge. That challenge concerns philosophers who deny some kinds of standard moral luck but accept others, and it consists in their having to come up with a principled reason for drawing the line where they do. The relevant difference task not only adds to the line-drawing challenge for these philosophers but also expands the challenge to include those who deny all kinds of standard moral luck.

⁷ Nagel (1979/2012) takes standard resultant, circumstantial, and constitutive luck to threaten agency. An interesting feature of luck with regard to moral principles is that it does not threaten agency in this way (a feature it shares with

Finally, I make some concluding remarks about the possibility to solve the relevant difference task by endorsing moral responsibility skepticism – the view that agents like us are never morally responsible – and summarize my arguments (section 7).

3. “The Usual Characterization of Moral Luck Is Faulty”

One strategy for standard moral luck deniers who want to avoid the relevant difference task is to resist the claim that moral luck occurs when factors beyond agents’ control affect moral responsibility. If moral luck deniers can successfully redefine moral luck so as to exclude moral principles as sources of it, they avoid the relevant difference task. I will consider two alternative ways of characterizing moral luck.⁸

First, deniers of standard moral luck might suggest that moral luck has to do with whether factors beyond an agent’s control affect what she is responsible *for* – whether she is or is not responsible for a death, for example. They might suggest, that is, that moral luck has to do with whether the *scope* of an agent’s moral responsibility is affected by factors beyond her control. The scope of moral responsibility can be distinguished from the *degree* of moral responsibility, which tells us *how* responsible an agent is – for instance, whether she is responsible to degree x or the lesser degree y . It can also be distinguished from the *valence* of moral responsibility, which specifies how an agent’s moral responsibility is *charged* – whether she is blameworthy or praiseworthy.⁹

The suggestion on the table aims to avoid the relevant difference task by excluding as sources of moral luck factors that affect only degree or valence of moral responsibility. To avoid the relevant difference task, standard moral luck deniers must, however, be more precise since it seems that moral principles *can* affect the scope of moral responsibility. Suppose a true moral principle says that agents who murder in Killer’s circumstances are blameworthy for doing so. This principle affects the scope of Killer’s moral responsibility since it is because the principle is true that she is blameworthy for murdering the target. If it is not true that agents who murder in Killer’s circumstances are blameworthy for doing so, Killer is not blameworthy for murder.

To get around this, deniers of standard moral luck could claim that moral principles have no causal effects and suggest the following definition:

what Story (2019) calls “interpersonal moral luck”). Note, however, that pointing to this feature as a relevant difference does not seem like a promising strategy because being a threat to agency is plausibly not always a feature of standard luck factors. It seems clear, for example, that standard luck factors do not have this feature if agent-causal libertarianism is true – in that case, there would still be constitutive, circumstantial, and resultant luck in peoples’ lives, but this fact would not, I take it, threaten to rob them of their agency.

⁸ Both are definitions that may come to mind in light of Nagel’s original definition of moral luck which says that “[w]here a significant aspect of what someone does depends on factors beyond his control, yet we continue to treat him in that respect as an object of moral judgment, it can be called moral luck” (1979/2012: 26).

⁹ Zimmerman (2002: 560–561) makes the distinction between scope and degree of moral responsibility, while to my knowledge Herdova and Kearns (2015: 364) introduce the valence terminology into the moral luck debate.

Scope. Moral luck occurs when, and only when, factors beyond agents' control *causally* affect the scope of their moral responsibility.

If moral principles have no causal effects, *Scope* rules them out as sources of moral luck. At the same time, *Scope* allows standard luck factors to be sources of moral luck since standard luck factors causally affect the scope of moral responsibility – for example, a bird intercepting Birdy's bullet causally affects what she is responsible for (attempted murder, but not a death). Standard moral luck deniers could argue, then, that there is no relevant difference task for them to deal with because when we pay closer attention to how the definition of moral luck should look, we see that even if moral principles affect moral responsibility, they do not give rise to moral luck.

For this strategy to be successful, deniers of standard moral luck must, to begin with, show that moral principles are causally inert, which is not obvious.¹⁰ Regardless of this issue, however, it seems to me that we should not accept *Scope*. This is because the definition implies that some views that should classify as accepting moral luck do not so classify.

Consider, first, a view on which factors beyond an agent's control affect her moral responsibility *tout court*. The notion of moral responsibility *tout court* is from Zimmerman (2002: 564–565), who maintains that agents can be morally responsible in virtue of something without being morally responsible *for* anything. In these cases, the agents are responsible *tout court*. By way of example, Zimmerman would say that an agent can be blameworthy to degree x in virtue of being such that she would attempt murder in certain circumstances, even if she has never attempted murder and thus is not responsible *for* that.

Now, consider two agents who do *not* get an offer to murder for money. One of them would, due to her upbringing and genetic makeup, accept such an offer if it were put to her, while the other would not accept it, due to her very different upbringing and genetic makeup. Suppose that the agent who would *not* accept the offer is less blameworthy *tout court* than the agent who *would* accept the offer, even though neither is responsible *for* taking the offer since neither got it in the first place. It seems that this is precisely the kind of phenomenon one wants to capture with the notion of moral luck. The agent who would *not* accept the offer seems to get off lightly solely because of luck – she seems to be morally lucky to be the way that she is if she is less blameworthy *tout court* than the agent who would accept the offer. Thus, a view on which this judgement is true should thereby classify as one accepting moral luck. Since the scope of these agents' moral

¹⁰ For instance, Tännsjö (2010: 52–58) seems open to the idea that moral facts – even general ones – can cause moral beliefs. In addition, Stringer (2018) argues that moral properties have causal powers and Oddie (2005: 181–210) that values do. Depending on what one takes moral principles to be more precisely, these views may imply that moral principles are causally efficacious.

responsibility is not affected, however, *Scope* implies that the view does not so classify. Hence, *Scope* is inadequate.

Second, consider a view on which the result of an action can affect degree of moral responsibility for that same action. Recall Birdy. Because a bird intercepts Birdy's bullet, her action does not result in a death. Suppose that, because of this, Birdy is less blameworthy for her action than she would have been, had that same action resulted in a death. If so, Birdy seems to get off lightly solely because of luck – she seems to be morally lucky not to kill the target. It thus seems that this too is precisely the kind of phenomenon one wants to capture with the notion of moral luck, and that a view on which this judgement is true thereby should classify as accepting moral luck. However, since the result of an action does not causally affect what action the agent is responsible for, *Scope* entails that the view does not so classify.

As I hope is clear, in endorsing these arguments, I am not endorsing the view that moral responsibility *tout court* can be affected by factors beyond agents' control, nor the view that the result of an action can affect degree of moral responsibility for that same action. Rather, my point is that these views should be classified as views that accept moral luck. *Scope* is inadequate, then, because it entails that they do not so classify.

A second possibility for deniers of standard moral luck who want to rule out moral principles as sources of moral luck by resisting the standard definition is to say that moral luck has to do with whether the fact in virtue of which an agent is responsible – her performing an action or causing some state of affairs, for instance – is beyond her control. At first, it might seem like this cannot get standard moral luck deniers very far since one may think that agents are responsible partly in virtue of moral facts. But standard moral luck deniers could insist, in order to rule out moral principles as sources of moral luck, that *moral luck* has to do with only the non-moral facts – that is, the purely descriptive facts – in virtue of which agents are responsible. They could suggest the following definition, then:

In-Virtue-Of. Moral luck occurs when, and only when, the non-moral fact in virtue of which an agent is responsible is beyond her control.

On this definition, moral luck occurs for instance if an agent is blameworthy in virtue of the fact that she attempted to kill someone and the fact that she attempted to kill someone was beyond her control. If *In-Virtue-Of* is adequate, moral principles do not give rise to moral luck and hence, there is no relevant difference task for deniers of standard moral luck to solve.

I believe, however, that *In-Virtue-Of* is inadequate because, just like *Scope*, it does not classify some views that should count as accepting moral luck as accepting it. But let me start by saying a few words about the relevant in-virtue-of relation. I have in mind a relation such that if an agent is

blameworthy in virtue of some fact, then that fact is *blameworthy-making* in the same way that wrong- and right-making features make actions right or wrong. Not just any fact that is relevant for blameworthiness is a part of the in-virtue-of base on this understanding, then, but only facts that are analogous to right- and wrong-making features – only facts that ground, or make it the case, that the agent is blameworthy.

To see how *In-Virtue-Of* encounters a classification problem, recall Killer and suppose that Killer's target never freely killed anybody. Then consider a view on which Killer is blameworthy for having killed her target, on which the target's never having freely killed anybody implies that Killer is more blameworthy for the killing than she would have been if the target had freely killed somebody, and on which the target's not having freely killed anybody is *not* a part of that in virtue of which Killer is blameworthy (perhaps this fact is more of a background condition according to the view under consideration). Because Killer does not control whether the target has freely killed somebody, Killer seems to be more blameworthy solely because of luck on this view – the view seems to imply that she is morally unlucky that the target never freely killed anybody. An adequate definition of moral luck should therefore imply that the view thereby classifies as a view that accepts moral luck. *In-Virtue-Of* entails that it does not so classify, however, since the factor beyond Killer's control affecting her moral responsibility – whether her target freely killed anybody – is not, on this view, a part of her responsibility's in-virtue-of base.

Note two things about this argument. First, the argument, just like those against *Scope*, does not commit me to the view under consideration, nor does it require that the view is correct. The argument only requires the view to be a not clearly implausible one that should classify as accepting moral luck. Second, it is not contradictory to endorse both the claim that a factor affects Killer's blameworthiness and the claim that the same factor is not a part of that in virtue of which Killer is blameworthy. If these claims were inconsistent, any factor that affects blameworthiness would have to be blameworthy-making – that is, would have to ground blameworthiness – but this is arguably not the case. The fact that Killer was neither struck dead by lightning nor abducted by aliens before taking a shot at her target affects her blameworthiness. Yet one need not, to avoid contradiction, accept that *Killer's not being struck dead by lightning before taking the shot* and *Killer's not being abducted by aliens before taking the shot* are a part of that which grounds Killer's blameworthiness for the killing.

In conclusion, neither attempt to rule out moral principles as sources of moral luck by redefining moral luck is successful.¹¹

¹¹ Another natural suggestion, noted by a referee, is to claim that moral luck occurs if and only if a factor affects an agent's moral responsibility *in a way* that she cannot control. This, one could argue, would rule out moral principles as sources of moral luck because, although moral principles affect moral responsibility, agents can control *how* moral principles affect their responsibility by acting, with control, one way or another. For instance, one could argue that an

4. “Neither Standard Luck Factors nor Moral Principles Affect Moral Responsibility”

So far, I have considered whether deniers of standard moral luck can avoid the relevant difference task by claiming that the correct definition of moral luck excludes moral principles as sources of it. Another way deniers of standard moral luck might try to avoid the relevant difference task is by claiming that moral principles do not, in fact, affect moral responsibility. In that case, they would not give rise to moral luck and hence, there would be no relevant difference task for standard moral luck deniers to solve. Below, I explore four ways of developing this line of reasoning and argue that none succeed.

Consider, first, an argument via a solution to the problem of moral luck put forward by Zimmerman (2002; 2015). Zimmerman (2002: 564–565) argues that agents are responsible in virtue of facts about their counterfactual free actions – facts about what actions they *would* freely perform in various circumstances. If this account is correct, one might argue that moral principles about the deontic status of actions do not affect moral responsibility because what agents *do* is irrelevant for determining moral responsibility.

Moral principles affect moral responsibility, however, even if agents are responsible in virtue of what they *would* do and not what they do. After all, principles about the deontic status of the actions they would freely perform affect moral responsibility on this picture, as do principles about whether those actions are such that agents are blameworthy or praiseworthy in virtue of them. Moreover, to the extent that an agent acts freely, her actions *are* relevant for determining her moral responsibility because they entail that she would freely perform them in those circumstances. That I freely φ in a set of circumstances entails, as Zimmerman himself notes (2015: 151), that I would freely φ if I were in those circumstances. Hence, one cannot argue by way of Zimmerman’s solution that moral principles do not affect moral responsibility.

Second, one might think that another account of moral responsibility that has been proposed in response to the problem of moral luck – one that grounds moral responsibility in character (e.g., Peels 2015; Rescher 1993; Richards 1986; Thomson 1989) – could do the job. According to this

agent controls how moral principles affect her responsibility by refraining from murder. I do not think this is a promising line of argument, however, because this definition, too, seems to encounter classification problems. For example, it encounters such problems with regard to the following classic circumstantial luck scenario: One German is sent by his employer to Argentina just before the rise of the Nazi regime and goes on to live a quiet life there, while another German is not sent to Argentina and goes on to partake in the Nazi regime’s crimes. Had the former not been sent to Argentina, he, too, would have partaken in the Nazi regime’s crimes (see Nagel, 1979/2012: 25–26, for the first mention of a case along these lines). A view on which the agents in this paradigmatic case are blameworthy to different degrees should thereby count, I take it, as a view that accepts moral luck. The present definition entails that it does not so classify, however, because both agents can control how their (not) being transferred affects their responsibility by acting, with control, one way or another.

account, agents are responsible in virtue of, and only in virtue of, features of their character.¹² What matters fundamentally for moral responsibility on this picture is not what agents do, but what they are like. If this is correct, one might argue that moral responsibility is not affected by moral principles about deontic status, nor by principles about what actions make agents blame- or praiseworthy, simply because how agents act is irrelevant for moral responsibility.

However, even if agents are indeed morally responsible in virtue of, and only in virtue of, features of their character, it is still the case that moral principles affect moral responsibility. Specifically, moral principles about what character traits make one blameworthy or praiseworthy affect moral responsibility on this picture. Take the trait of being greedy, as Sneezy is. The fact that being greedy is a character trait that agents are blameworthy in virtue of can affect Sneezy's moral responsibility; it is because greediness is a character trait that makes agents blameworthy that Sneezy is blameworthy in virtue of it. One cannot argue, then, that moral principles do not affect moral responsibility by appealing to this account of what agents are responsible in virtue of.

A third possible argument for the conclusion that moral principles do not affect moral responsibility is that with regard to moral responsibility, what matters is not the moral principles *per se* but what agents *believe* about them. Suppose, for example, that someone wholeheartedly believes that donating to a certain charity is right and has done all that is required to non-culpably believe so. One might argue that even if it is in fact wrong to donate to that specific charity because, say, the money is used to fund trafficking rather than alleviate poverty, the agent is not more blameworthy for donating than she would have been, had it been right to do so. If she is not, deontic status *per se* does not affect moral responsibility, and the same could be said for moral principles about what kind of actions make agents blame- or praiseworthy.

Nevertheless, even if the specific principles I have used so far to make my case do not affect moral responsibility if beliefs are what matters for it, other principles still do. Consider the control principle itself, for instance. Clearly, regardless of one's beliefs about the control principle, one's moral responsibility can be affected by whether one can indeed be responsible for something only to the extent that it depends on factors within one's control. Take an assassin who shoots at her target, who believes that she will be blameworthy for a death if her bullet hits home, and who thanks to factors beyond her control does hit home. Regardless of what she believes, she is not blameworthy for a death if the control principle is true, and it is because the control principle is true that she is not.

¹² I follow Hartman (2020: 106) in describing the core of these views like this even though some of them, as Hartman notes, differ in important respects.

Furthermore, moral responsibility is on this picture affected by moral principles about the relation between beliefs and moral responsibility – for example, by whether a correct moral principle says that a certain belief is such that it makes agents blameworthy for an action to degree x rather than y , or blameworthy rather than praiseworthy for it. Suppose I perform an action, believing that a correct moral principle says that agents who perform an action like that in my circumstances are blameworthy for doing so. Whether a correct moral principle says that agents are blameworthy for performing actions they believe they will be blameworthy for affects my responsibility: if a correct moral principle says that it is the case, then I am blameworthy for what I have done, and it is because it says so that I am. Moral principles affect moral responsibility even if beliefs are what matters for it, then.

Fourth, and finally, one might argue that moral principles appear to affect moral responsibility only because we have been trading with too wide a notion of affecting. Once that notion has been properly specified it becomes clear, one could argue, that moral principles do not affect moral responsibility – and hence, that standard moral luck deniers do not face the relevant difference task. One might argue, for instance, that for something x to affect something y in the sense at issue in the standard characterization of moral luck, it must be the case that y holds in virtue of x (that is, as mentioned in section 3, that x grounds y or makes it the case that y holds). And since an agent is not responsible for an action in virtue of moral principles, but in virtue of purely non-moral facts such as her having killed someone (or so the argument would go), moral principles do not affect moral responsibility.

By pursuing this line of argument, however, one is defining away the problem of moral luck. On the critic's notion of affecting, standard luck factors (such as whether a bird intercepts an assassin's bullet) give rise to moral luck only if they ground moral responsibility, since it is only then that they would affect moral responsibility on that understanding. But standard luck factors are not, it would seem, such that they ground moral responsibility. Even if Killer is morally responsible in virtue of purely non-moral facts, it is plausibly not the non-moral fact that no bird intercepted her bullet that grounds her responsibility – it seems that such a grounding fact must have something to do with *her*, like her having killed her target. Likewise, even if Birdy is morally responsible for an attempt to kill in virtue of non-moral facts only, what grounds her responsibility is surely not the non-moral fact that she didn't sneeze at a crucial moment. And the same goes for Sneezy: even if Sneezy is morally responsible for being greedy and violent in virtue of non-moral facts only, it seems it is not the non-moral fact that her upbringing and genetic makeup resulted in those traits that grounds her responsibility. The proposed notion of affecting thus excludes not only the possibility that moral principles give rise to moral luck, but also the possibility that standard

luck factors do so. Hence, by defining ‘affect’ this way, the standard moral luck denier would rule out the existence of the phenomenon whose existence is the very topic of the conversation. This is not an acceptable move in trying to avoid the relevant difference task.

It seems hard to deny, then, that moral principles affect moral responsibility.¹³ If deniers of standard moral luck cannot find some other way to avoid the relevant difference task, they must show us how moral principles relevantly differ from standard luck factors to maintain their position. In the following sections, I explore two natural candidates for relevant differences.

5. “Moral Principles Are Necessary, Standard Luck Factors Are Not”

Looking for a relevant difference between moral principles and standard luck factors, the first thing that comes to mind is perhaps that it is in the modality. While standard luck factors are contingent, moral principles, one might contend, are necessary. Hence, deniers of standard moral luck could hold that while necessary factors beyond agents’ control can affect moral responsibility, contingent factors cannot.

Notice, first, that this reasoning is, at least in one important respect, somewhat surprising. The motivation for denying standard moral luck is that its existence conflicts with the idea that agents are responsible for something only to the extent that it depends on factors within their control. Therefore, one would not expect deniers of standard moral luck to argue that moral responsibility can be affected by factors so utterly beyond agents’ control – factors that it is metaphysically impossible for them to make false. Something does not seem to sit quite right with that combination of ideas.

A second, more pressing, worry for those who would argue that modality is the relevant difference, is that doing so commits them to the claim that whether moral principles can affect moral responsibility *hinges on* whether they are necessary. Suppose a correct principle says that agents who murder in Killer’s circumstances are blameworthy for doing so. Now if I, in the way suggested above, appeal to modality as the relevant difference between moral principles and standard luck factors, I must say that this principle can affect Killer’s moral responsibility if it is true in all possible worlds. But if it is not – if the principle is contingently true – I must say that it cannot affect moral responsibility. Plausibly, however, whether a correct moral principle affects moral responsibility does not *depend* on whether the principle is true in all possible worlds. If Killer actually murders her target and a principle that is true in the actual world implies that she is

¹³ It has been argued that moral principles (as opposed to particular moral facts) are explanatorily inert (Berker 2019). Although I use the example of moral principles, one could easily recast the challenge in terms of particular moral facts, or even moral properties. Thus, deniers of moral luck cannot get off the hook by saying that moral principles (as opposed to particular moral facts or moral properties) do not affect moral responsibility because they are explanatorily inert.

blameworthy for doing so, then surely whether this affects her responsibility does not depend on whether the principle is true in all other possible worlds as well.

It further seems implausible, as we have seen in section 4, that correct moral principles do not affect moral responsibility. If it is true in the actual world that agents who murder in circumstances like Killer's are blameworthy for doing so, it is very difficult to see how that could fail to affect the responsibility of someone who murders in those circumstances. But that is precisely what standard moral luck deniers that point to the modality difference in this way will have to hold, if it turns out that moral principles are contingent rather than necessary, as some do argue (e.g., Hattiangadi 2018; Rosen 2020; Tännsjö 2010: 48–49). To put the point differently: this way of appealing to modality as the relevant difference commits deniers of standard moral luck to a second implausible claim; namely, that contingent moral principles do not affect moral responsibility.¹⁴

Finally, and most importantly, deniers of standard moral luck must explain why the difference in modality is relevant. And discerning why necessity would make moral principles *relevantly* different from contingent standard luck factors is difficult. To see this, recall Birdy the unsuccessful assassin and suppose you have, up until now, denied standard moral luck. Imagine also that you recently came to believe that God exists and is necessarily almighty, that God's will is necessary, and that God willed a bird to intercept Birdy's bullet. Given your new beliefs about God, you should now believe that a bird's intercepting Birdy's bullet was necessary – you should now believe that the luck factor in Birdy's scenario is necessary rather than contingent. Because of this, if you have appealed to modality by saying that necessary factors beyond agents' control can affect moral responsibility while contingent factors cannot, you would now have to change your mind about the existence of standard moral luck. You would now have to hold that whether a bird intercepts Birdy's bullet *can* affect her moral responsibility and thus, that standard moral luck exists.

That seems like a strange result, however, since it is still the case that Birdy gets off lightly solely because of luck if she is judged less blameworthy than Killer. To the extent that you, like standard moral luck deniers, deny moral luck because you are concerned with keeping the scope, degree or valence of moral responsibility luck-free, it seems that you ought not to change your moral responsibility judgment simply because you come to believe that the factor beyond Birdy's control is necessary, rather than contingent. This indicates that a difference in modality between factors is irrelevant.

One may still think that modality, while not directly relevant, nevertheless somehow plays an important indirect role. Perhaps one believes, for instance, that since moral principles are the same for all agents, it would not be unfair for them to affect responsibility while it would be unfair for

¹⁴ See also Dreier (2019: 1404–1407), who argues that moral principles give rise to moral luck if they are contingent.

standard luck factors to do so, since they are not the same for all agents. Modality matters in this case because the reasoning is especially plausible if moral principles are necessary, since they would then be the same for all agents across possible worlds.

However, since this argument relies on moral principles's being necessary, this account of a relevant difference inherits the commitment to the problematic claims that pointing directly to modality in the way discussed above involves. Moreover, the supposed unfairness of standard moral luck seems to derive from agents' lack of control. When we look at comparisons that are used to illustrate moral luck (comparisons like *Results*, *Circumstances*, and *Constitution*), it may seem unfair to judge the agents blameworthy to different degrees and for that reason, one may want to deny moral luck. But the judgement that the agents differ in their blameworthiness seems unfair *because* they lack control over what makes the difference between them. If the agents controlled the difference between them, it would not be unfair to judge them blameworthy to different degrees. The lack of control seems primary, then, and the unfairness of moral luck a result of the lack of control. Whether the moral principles are the same for everyone is not relevant, then, since even if they are, they affect moral responsibility while being beyond agents' control.

On a related note, someone might suspect that comparisons between agents, which are often used to explicate moral luck, are essential for it. If so, moral luck deniers could argue that standard luck factors and moral principles relevantly differ in that the former, but not the latter, can be variables in possible comparisons between agents. Modality matters indirectly in that case because the reasoning presupposes that moral principles are necessary – only then would it be impossible for them to be such variables.

I do not think, however, that comparisons are essential to moral luck. It seems to me, rather, that their function is illustrative. When one wants to show that it looks like moral responsibility depends in more ways than previously recognized on factors beyond our control, comparisons are a convenient way to do so. By using a comparison like *Results*, for instance, one can make it salient that whether assassins succeed in what they try to do depends on things like the absence of birds in their vicinity. Because the comparison makes this salient, it also makes it clear that we accept moral luck when we accept, as it seems we do in our everyday moral judgments, that agents who succeed in their efforts to murder are more blameworthy than agents whose efforts fail. Of course, if a factor is necessary, it cannot figure in a comparison. That does not show, however, that such factors do not give rise to moral luck, but only that it is impossible to illustrate that we seem to endorse that they do by way of a comparison. The idea that comparisons serve merely an illustrative purpose is further corroborated by the fact that Nagel, who delineated the moral luck problem in its present form, introduces comparisons not to show that moral luck arises only from factors that

can figure in possible comparisons, but to illustrate that moral responsibility always depends on factors beyond agents' control (1979/2012: 25–26).

The relevant difference task cannot, it seems, be solved by a direct or indirect appeal to the difference in modality between standard luck factors and moral principles. When it comes to moral luck, what matters is whether agents have control over the factors that affect moral responsibility and when they do not, it is difficult to see why it would be pertinent that the factor beyond their control is necessary rather than contingent.

6. “Standard Luck Factors, but Not Moral Principles, Are Matters of Luck”

Another suggestion close at hand when trying to find a relevant difference between standard luck factors and moral principles is that while the truth of moral principles is not a matter of luck, standard luck factors are matters of luck. One might believe, for example, that only things that could easily have failed to occur are lucky (Coffman 2014: 494; Peels 2015: 77; Pritchard 2005: 128), or that only things that come to be by pure chance are (Rescher 2014: 621). Arguably, if that is so, the truth of moral principles is not a matter of luck since their truth could not easily have failed to obtain and did not come to be by pure chance. Thus, deniers of standard moral luck could claim that while moral principles can affect moral responsibility, since they are *not* matters of luck, standard luck factors cannot, because they *are* matters of luck.

Let us accept, for the sake of argument, that all standard luck factors are matters of luck and that moral principles are not. Even so, it is not clear why that would be *relevant* since the problem at hand is about luck only to the extent that the term ‘luck’ designates lack of control. As previously mentioned, the contradiction that the problem of moral luck consists in arises because common sense seems to have it, on the one hand, that agents are morally responsible for something only to the extent that it depends on factors within their control and, on the other hand, that agents differ in moral responsibility even though the difference between them is due to factors over which they lack control. So, as Hartman puts it, “the same debate could continue on seamlessly without recourse to the word ‘luck’, because the paradox arises on account of the tension between ubiquitous lack of control and our conception of morality” (2017: 29).¹⁵

Some may find this too quick a dismissal and insist that it is crucial for fully understanding moral luck that lack of control is necessary but not sufficient for luck.¹⁶ My arguments do not depend on my denying this, however. Even if there are important issues pertaining to moral luck that can be dealt with only by using a narrower notion of luck, one important part of the debate concerns tensions in our beliefs about moral responsibility and mere lack of control. That is the

¹⁵ Hartman (2017: 23–31) offers an in-depth discussion of the issue.

¹⁶ See, for example, Levy (2019).

part that I, like many others, am concerned with – and there are no particular reasons to think that this part of the debate is insubstantial.

Of course, even if there is no relevant difference between standard luck factors and moral principles to point to concerning luck, one could deny that moral principles are beyond agents' control. In that case, moral principles are ruled out as sources of moral luck even on the lack of control conception of luck. One way to go about this line of argument is to say that an agent has control over p if she has an action available to her such that if she performs it, p is true. On this notion of control (let us call it “ensure-that control”), agents control moral principles if they are necessary. After all, if moral principles are necessary, an agent always has an action available to her such that if she performs it, the moral principle is true.

Although it is not entirely clear what notion of control is at issue in the moral luck debate, it seems that it is not ensure-that control. To appreciate this, take Killer and Birdy and suppose that Killer is causally determined to murder her target, that Birdy is causally determined *not* to murder her target, and that both assassins are morally responsible despite being causally determined. If ensure-that control were at issue in the moral luck debate, those who usually deny that Killer is more blameworthy than Birdy because Killer caused a death should now change their minds. Since Killer has ensure-that control over killing her target, and Birdy has ensure-that control over not killing her target, the difference between the two assassins is not due to factors beyond their ensure-that control. Thus, there should be no problem with judging them blameworthy to different degrees. I doubt deniers of standard moral luck will be convinced that there is no problem with this judgement, however. It still does not seem like the difference between the agents is up to them in a sense that is crucial for judging them blameworthy to different degrees. Some may think that what is needed is that the agents control *whether* they murder their target, not *that* they murder and do not murder their target, respectively. Others will deny this because they believe that the ability to do otherwise is irrelevant to moral responsibility.¹⁷ Whatever the notion of control at issue in the moral luck debate is more precisely, it is clearly not ensure-that control.

It is a benign assumption, then, that agents do not control moral principles.¹⁸ This seems so even if there are some notions of control according to which agents do control moral principles. For any plausible account of the notion of control *at issue in the moral luck debate* will surely rule out that agents have control over such things as moral principles. Because this puts moral principles in

¹⁷ See the debate over so-called Frankfurt cases, initiated by Frankfurt (1969).

¹⁸ Another way to deny that moral principles are beyond agents' control might be to embrace a kind of individual moral relativism – one where the truth or falsity of moral propositions is relative to individuals. If that is the case, then perhaps moral principles are under agents' control in some sense. It seems unlikely, however, that deniers of standard moral luck would want to base their defense on this position.

the category of things over which agents lack control, together with standard luck factors, there is no relevant difference between the two to point to concerning luck or control.

7. Concluding Remarks

In closing, a few words are in order about a possibility that is bound to have been in the back of some readers' minds throughout, namely, that one can solve the relevant difference task by endorsing moral responsibility skepticism, that is, the view that agents like us are never morally responsible.

I take it that this view does not allow one to categorically deny the existence of moral luck, because whether moral responsibility skepticism is true seems to affect moral responsibility – for instance, its truth rules out that Killer is morally responsible for a murder. Endorsing moral responsibility skepticism does, however, provide one with a relevant difference that allows for the categorical denial of *standard* moral luck, namely, that luck factors that can affect moral responsibility if responsibility skepticism is true give rise to moral luck, while luck factors that cannot do not give rise to moral luck. If moral responsibility skepticism is true, moral responsibility would not be affected by standard luck factors – moral responsibility would not be affected, for instance, by whether agents take up offers to murder for money, whether they get an opportunity to go through with their plans, and whether they succeed in them. After all, no agent is ever responsible for anything if moral responsibility skepticism is true.

Because standard moral luck deniers' motivation for denying moral luck in the first place is that they believe it conflicts with the control principle, one might suspect that it is an inconsistent move on their part to endorse moral responsibility skepticism. Just as it conflicts with the control principle that standard luck factors affect moral responsibility, so too it may seem – since agents do not control whether moral responsibility skepticism holds – that moral responsibility skepticism's affecting moral responsibility conflicts with the control principle. There is no such conflict, however. The control principle states that an agent is morally responsible for something only to the extent that it depends on factors within her control. But of course, if no one is ever responsible for anything at all, no agent will ever be responsible for something to an extent that goes beyond her control. Hence, although the truth of moral responsibility skepticism affects moral responsibility, it does not conflict with the control principle. By appealing to moral responsibility skepticism, deniers of standard moral luck can thus solve the relevant difference task. This can get

some deniers of standard moral luck off the hook, but it will surely be an unattractive option for most of them since nothing indicates that they are willing to give up responsible agency.¹⁹

To sum up, I have in this paper presented a challenge for deniers of standard moral luck. I have argued that, because it seems that some factors give rise to moral luck, they face the dialectical burden of either pointing to a relevant difference between these factors and the factors that according to them do not give rise to moral luck, or showing how they can avoid having to do so. After introducing the challenge, I considered the prospects of solving it, using moral principles as an example. I first looked at whether standard moral luck deniers can avoid having to solve the relevant difference task by claiming either that the correct definition of moral luck rules out moral principles as sources of it, or that moral principles do not, in fact, affect moral responsibility. I argued that both strategies are unsuccessful. Second, I considered and criticized two natural candidates for a relevant difference: that moral principles, but not standard luck factors, are necessary, and that standard luck factors, but not moral principles, are matters of luck. I argued that neither amounts to a relevant difference and, in addition, that pointing to the former commits deniers of standard moral luck to some surprising and implausible claims.

Naturally, my arguments against these suggestions do not demonstrate that there is no relevant difference, nor that there is no way for standard moral luck deniers to show that they need not point to a relevant difference. But I hope to have shown at least that there is some dialectical slack for standard moral luck deniers to pick up and that, if they cannot rise to the occasion, they are forced to accept standard moral luck unless they are willing to become moral responsibility skeptics.

Acknowledgments

I am very grateful to my Ph.D. supervisors, Erik Carlson and Jens Johansson, for their valuable guidance and meticulous feedback. Many thanks to the latter also for making the remark that prompted me to write this paper; namely, that moral principles seem to give rise to moral luck. For constructive questions and discussion, I would like to thank the participants at the Higher Seminar in Practical Philosophy at Uppsala University, the 2022 Swedish Congress of Philosophy, and the

¹⁹ Hartman defines moral luck as “occurring when factors beyond an agent’s control partially determine her positive praiseworthiness or blameworthiness”, where the term ‘positive’ is “meant to rule out the idea that moral luck is responsibility-undermining luck” (2017: 2). On this definition, the truth of responsibility skepticism does not give rise to moral luck. It is not clear to me, however, that we should endorse Hartman’s definition. Consider agents like Birdy. If Birdy, due to a bird’s intercepting her bullet, is not blameworthy for a murder and less blameworthy than she would have been if no bird had intercepted her bullet, we should say, I take it, that she is subject to moral luck. But this luck *is* responsibility-undermining – it is luck that implies that Birdy is not responsible for a murder, and overall blameworthy to a lesser degree than she would have been if a bird had not intercepted her bullet. Note, however, that also those who find Hartman’s definition attractive can accept the main claim that I make concerning responsibility skepticism; namely, that although endorsing this view deals with the relevant difference task, it is likely an unattractive option for most deniers of standard moral luck.

26th Oxford Graduate Philosophy Conference. Finally, thanks to Roger Crisp, who served as a commentator on the paper in Oxford, and two anonymous referees at *Ergo* for their helpful comments.

References

- Berker, Selim. (2019). The Explanatory Ambitions of Moral Principles. *Nous*, 53(4), 904–936.
<https://doi.org/10.1111/nous.12246>
- Coffman, E. J. (2014). Strokes of Luck. *Metaphilosophy*, 45(4–5), 477–508.
<https://doi.org/10.1111/meta.12097>
- Dreier, Jamie. (2019). Is There a Supervenience Problem for Robust Moral Realism? *Philosophical Studies*, 176(6), 1391–1408. <https://doi.org/10.1007/s11098-019-01244-w>
- Enoch, David and Andrei Marmor (2007). The Case Against Moral Luck. *Law and Philosophy*, 26(4), 405–436. <https://doi.org/10.1007/s10982-006-9001-3>
- Frankfurt, Harry G. (1969). Alternate Possibilities and Moral Responsibility. *The Journal of Philosophy*, 66(23), 829–839. <https://doi.org/10.2307/2023833>
- Hanna, Nathan. (2014). Moral Luck Defended. *Nous*, 48(4), 683–698.
<https://doi.org/10.1111/j.1468-0068.2012.00869.x>
- Hartman, Robert J. (2017). *In Defense of Moral Luck*. Routledge, Taylor & Francis Group.
- Hartman, Robert J. (2020). Against the Character Solution to the Problem of Moral Luck. *Australasian Journal of Philosophy*, 98(1), 105–118.
<https://doi.org/10.1080/00048402.2019.1617755>
- Hattiangadi, Anandi. (2018). Moral Supervenience. *Canadian Journal of Philosophy*, 48(3–4), 592–615. <https://doi.org/10.1080/00455091.2018.1436034>
- Herdova, Marcela and Stephen Kearns. (2015). Get Lucky: Situationism and Circumstantial Moral Luck. *Philosophical Explorations*, 18(3), 362–377.
<https://doi.org/10.1080/13869795.2015.1026923>
- Jensen, Henning. (1984). Morality and Luck. *Philosophy*, 59(229), 323–330.
<https://doi.org/10.1017/S0031819100069916>
- Khoury, Andrew C. (2018). The Objects of Moral Responsibility. *Philosophical Studies*, 175(6), 1357–1381. <https://doi.org/10.1007/s11098-017-0914-5>

- Levy, Neil. (2011). *Hard Luck: How Luck Undermines Free Will and Moral Responsibility*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199601387.001.0001>
- Levy, Neil. (2019). Putting the Luck Back Into Moral Luck. *Midwest Studies In Philosophy*, 43(1), 59–74. <https://doi.org/10.1111/misp.12104>
- Lockhart, Jennifer Ryan and Thomas Lockhart (2018). Moral Luck and the Possibility of Agential Disjunctivism. *European Journal of Philosophy*, 26(1), 308–332. <https://doi.org/10.1111/ejop.12226>
- Nagel, Thomas. (2012). Moral Luck. In *Mortal Questions* (24–38). Cambridge University Press. <https://doi.org/10.1017/CBO9781107341050.005> (Original work published 1979)
- Nelkin, Dana Kay. (2019). Thinking Outside the (Traditional) Boxes of Moral Luck. *Midwest Studies In Philosophy*, 43(1), 7–23. <https://doi.org/10.1111/misp.12101>
- Nelkin, Dana Kay. (2021). Moral Luck. In Edward N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2021). Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/sum2021/entries/moral-luck/>
- Oddie, Graham. (2005). *Value, Reality, and Desire*. Oxford University Press. <https://doi.org/10.1093/0199273413.003.0007>
- Peels, Rik. (2015). A Modal Solution to the Problem of Moral Luck. *American Philosophical Quarterly*, 52(1), 73–87.
- Pritchard, Duncan. (2005). *Epistemic Luck*. Oxford University Press. <https://doi.org/10.1093/019928038X.001.0001>
- Rescher, N. (1993). Moral Luck. In Daniel Statman (Ed.), *Moral Luck* (141–166). State Univ. of New York Press.
- Rescher, Nicholas. (2014). The Machinations of Luck. *Metaphilosophy*, 45(4–5), 620–626. <https://doi.org/10.1111/meta.12101>
- Richards, Norvin. (1986). Luck and Desert. *Mind*, 95(378), 198–209. <https://doi.org/10.1093/mind/XCV.378.198>
- Rivera-López, Eduardo. (2016). How to Reject Resultant Moral Luck Alone. *Journal of Value Inquiry*, 50(2), 415–423. <https://doi.org/10.1007/s10790-015-9531-8>
- Rosebury, Brian. (1995). Moral Responsibility and “Moral Luck”. *The Philosophical Review*, 104(4), 499–524. <https://doi.org/10.2307/2185815>

- Rosen, Gideon. (2020). What is Normative Necessity? In Mircea Dumitru (Ed.) *Metaphysics, Meaning, and Modality: Themes from Kit Fine* (205–233). Oxford University Press.
<https://doi.org/10.1093/oso/9780199652624.003.0011>
- Story, Daniel. (2019). Interpersonal Moral Luck and Normative Entanglement. *Ergo* 6(21), 601–616. <https://doi.org/10.3998/ergo.12405314.0006.021>
- Stringer, Ryan. (2018). Realist Ethical Naturalism for Ethical Non-Naturalists. *Philosophical Studies*, 175(2), 339–362. <https://doi.org/10.1007/s11098-017-0870-0>
- Sverdlik, Steven. (1988). Crime and Moral Luck. *American Philosophical Quarterly*, 25(1), 79–86.
- Tännsjö, Torbjörn. (2010). *From Reasons to Norms: On the Basic Question in Ethics*. Springer.
<https://doi.org/10.1007/978-90-481-3285-0>
- Thomson, Judith Jarvis. (1989). Morality and Bad Luck. *Metaphilosophy*, 20(3–4), 203–221.
<https://doi.org/10.1111/j.1467-9973.1989.tb00423.x>
- Zimmerman, Michael J. (2002). Taking Luck Seriously. *The Journal of Philosophy*, 99(11), 553–576.
<https://doi.org/10.2307/3655750>
- Zimmerman, Michael J. (2015). Moral Luck Reexamined. In David Shoemaker (Ed.), *Oxford Studies in Agency and Responsibility: Volume 3* (136–159). Oxford University Press.
<https://doi.org/10.1093/acprof:oso/9780198744832.003.0008>