On a debunking explanation of moral luck

Tom Porter

Draft of March 2013. Please don't quote or cite without permission.

If appropriate moral assessment of a person can vary according to factors that are, from the point of view of that person, a matter of luck, then there can be *moral luck*. So, for example, there can be moral luck if I merit less severe moral assessment because I fail to fulfil my murderous intentions only thanks to an unanticipated lightning bolt that beats me to my intended victim. This particular form of moral luck is known as *resultant* moral luck, arising, as it does, from luck in the way that things turn out.[1] In this paper, I focus on resultant moral luck. There can be other kinds, but I don't talk about them here.

Something that suggests that there can be moral luck is our tendency to judge murderers more harshly than those who attempt murder but fail, for instance, and to judge those whose negligence results in disaster more harshly than those whose does not. On the other hand, something that speaks against the possibility of moral luck is the intuitive force of the 'control principle', according to which a person cannot be judged morally for what she did not control. So, both believers in and sceptics about moral luck have fairly powerful intuitions about specifically resultant moral luck on their sides.

As a result, the details of the debate between them must involve more than mere appeal to intuitions. And so it does. Believers in moral luck try to emphasise the centrality of luck, including moral luck, in our thinking

[1] The term is due to Zimmerman (1987: 376), although the form itself was originally distinguished by Nagel (1979).

about the good life, for example, and suggest that the control principle is too crudely formulated. Sceptics about moral luck, meanwhile, argue that the control principle is at the heart of our practices of moral assessment, and deny, therefore, the possibility of moral luck. And they attempt in various ways to debunk the appearance of acceptance of resultant moral luck in comparisons between successful and merely would-be murderers, for example.

In this paper, I focus on one such attempt to debunk the apparently moral luck-supporting intuitions that has been favoured by a number of sceptics about moral luck. The debunking explanation runs as follows. In cases of resultant moral luck such as the one I described above, in which an agent's morally improper intentions or reasoning produce no bad consequences only thanks to forces beyond her control, our intuition that she is to be judged less severely than an agent who was not so 'lucky' need not be interpreted as indicating a commitment to moral luck. It need be interpreted only as indicating a commitment to the appropriateness of worse *experiences* for the 'unlucky' agent than for the 'lucky' agent. But the justifications that can be offered for this need not be that the 'lucky' agent is appropriately to be morally judged more leniently. These justifications are, therefore, consistent with the view that in fact the 'lucky' agent and the 'unlucky' agent ought to be equally severely morally judged, as sceptics about moral luck believe.

In response, I consider a variation on one of the types of case for which the explanation is offered. The explanation in this type of case draws on the appropriateness of 'agent-regret', which is a kind of painful appreciation that one ought to feel of harms to others that are brought about through one's agency. I argue that this explanation doesn't work in the variation of the case that I describe, because even though we continue to have the

intuition that an 'unlucky' agent is to be judged more severely than the 'lucky' agent, there are no harms around to be a focus for agent-regret. Thus, the moral luck sceptics' proposed debunking of the seeming support for resultant moral luck that is provided by our intuitions about cases fails.

I begin, in section 1 below, by introducing a terminological framework for my discussion. There are three distinct families of moral concepts at work in debates about moral luck, but many of the terms in which the debate is conducted have senses from more than one family. So, for the sake of clarity, I stipulatively assign just one sense to each the terms that I'll be using. Then, in the second section, I introduce the puzzle of moral luck in more detail, and describe three possible views that one might adopt in response. In the third section, I set out the debunking explanation that sceptics about moral luck have offered of apparently moral luck-supporting intuitions about certain types of case, before describing, in the fourth section, a variation on one such type of case that continues to prompt the intuitions but in such a way that the debunking explanation fails. This variation constitutes, therefore, a serious challenge to the sceptical debunking strategy. In the fifth and final section, I consider and reject four objections.

1. *Frameworks*

A good place to start is with two cases described by Thomas Nagel in his seminal paper on moral luck. He writes

> If someone has had too much to drink and his car swerves on to the sidewalk, he can count himself morally lucky if there are no pedestrians in its path. If there were, he would

> be to blame for their deaths … But if he hurts no one, although his recklessness is exactly the same, he is guilty of a far less serious legal offence and will certainly reproach himself and be reproached by others much less severely. (1979: 29)

He continues:

> [T]he penalty for attempted murder is less than that for successful murder—however similar the intentions and motives of the assailant may be in the two cases. His degree of culpability can depend, it would seem, on whether the victim happened to be wearing a bullet-proof vest, or whether a bird flew into the path of the bullet—matters beyond his control. (1979: 29)

Many people make the judgments about these cases that Nagel describes. They have a sense that it is appropriate to reproach the lucky driver, who hits no one, much less severely than it would be appropriate to reproach him if he had hit someone. They also have a sense that the would-be murderer who kills no one is less culpable than the would-be murderer whose attempt to kill succeeds.

It's hard to say what the implications are of these views, however, without some clarification of the distinctions and connections between reproach, culpability, and blame. So, let me begin by distinguishing three related but nevertheless different families of concepts. My aim here is not to give a full description of the content of each concept—I assume that they will be familiar enough for that to be unnecessary—but to place them in

relation to each other in a way that offers us a useful and plausible framework for subsequent discussion. The framework that I offer will involve a certain amount of stipulation, because some of the words for the concepts that I distinguish also have senses that refer to others, and I shall disregard these other senses for the sake of clarity. If you take these other senses to be primary, you can map your favoured terminology onto the framework that I describe.

The first family, then, includes just one concept: *blameworthiness*. I shall say that a person is blameworthy for some action if, and only if, the action constitutes a wrongdoing and is attributable to her agency (cf. Enoch and Marmor 2006: 412). Sometimes, it might look as if a person is blameworthy for some wrongdoing, but then turn out that she isn't blameworthy after all, because the wrongdoing is not attributable to her agency (because her brain had been taken over by aliens, for example, or because she was blown by an unforeseeable giant gust of wind). In that case, we may relocate the blameworthiness (as we would in the aliens example), or we may decide that there was no wrongdoing in the first place (as we would in the wind example).

The second family of concepts includes *remorse*, *blame*, *resentment*, and *indignation*. Remorse is the appropriate disapproving sentimental reaction to one's own blameworthiness, and only to one's own blameworthiness. (One is prompted to remorse by one's conscience.) Greater remorse is appropriate to the extent that one is more blameworthy, or blameworthy for a more serious wrongdoing. Blame, I shall say, is the appropriate disapproving sentimental reaction to another's blameworthiness, and only to her blameworthiness (cf. Parfit 2011: 154). (It is *not* the public *expression* of that reaction: here, I stipulatively exclude that sense of the word 'blame' and the associated sense of the word 'blameworthy'.) Just as one should feel

more remorse to the extent that one is more blameworthy, so one is to be blamed more to the extent that one is more blameworthy. Resentment is a species of blame, and appropriate just when the wrongdoing for which another is blameworthy constitutes a wronging *of oneself*. Indignation is also a species of blame, and appropriate just when the wrongdoing for which another is blameworthy constitutes a wronging of some third party.

The third family of concepts includes public expressions of blame, and contains *censure*, *criticism*, and *reproach*. That these are distinct from what they publicly express can be seen from the fact that we don't always think that it's appropriate to censure even when we are confident that remorse and blame are appropriate. Sometimes we say, for example, that someone has suffered enough at the hands of her conscience, so that to reproach or criticise her would be inappropriate, even though we do blame her.

To summarise: the three families of concepts relate respectively to the attribution of wrongdoing, the appropriate sentimental reactions to wrongdoing, and the expressions of those reactions. To repeat: since, in my choice of the words that pick out the relevant concepts, I mean to exclude other perfectly acceptable senses of those words, my characterisation of the framework is to some extent stipulative. But the distinction between the three families and the relations between them are, I trust, clear enough.

2. *The puzzle of moral luck*

I said that many people make the judgments that Nagel describes in the passages that I quoted above. We're now in a position to understand the implications of that.

Suppose that the judgments that Nagel describes are redescribed in terms of the framework that I've just outlined, as follows. First, it is appropriate to *blame* the lucky driver, who hits no one, much less severely than it would be appropriate to blame him if he had hit someone. That is to say that the lucky driver is *less blameworthy* than he would have been if he had hit someone, and should feel *less remorse*. Second, the unsuccessful would-be murderer is also *less blameworthy* than he would have been had he succeeded in his murderous attempt, and so is also less to *blame* and is to be *reproached* less than he ought to have been had he succeeded.

This is how most philosophers have interpreted what Nagel says about the two cases, since this interpretation gives rise to the puzzle that Nagel is interested in describing. For, if you take the intuitive judgments at face value, then you believe in the possibility of *moral luck*. That is: you believe that the degree to which someone can be blameworthy and should feel remorse for some wrongdoing may be at least in part dependent upon luck. In particular, you believe in the possibility of *resultant* moral luck—the possibility of differences in blameworthiness that arise as a result of luck in the way that things turn out.

There are two positions available to a believer in the possibility of resultant moral luck.[2] One is

> *The concessive view*: that an act's blameworthiness cannot depend *entirely* on luck. But when two acts are blameworthy in some way that does not depend on luck, one of these acts may be *more* blameworthy in some way that *does* depend on luck.

[2] The descriptions that I offer here of the positions, but not their names, are taken from Parfit (2011: 155–6).

The other is

> *The extreme view*: that an act's blameworthiness might depend entirely on luck.

Both views are compatible with the judgments that Nagel describes. The difference between them can be illustrated as follows. Extremists could accept that the lucky driver and the unsuccessful would-be murderer were *not blameworthy at all*, even though each would have been blameworthy had his actions resulted in a person's death.[3] Since the difference between hitting someone and not hitting someone, from the driver's perspective, and the difference between being successful and unsuccessful, from the would-be murderer's perspective, are matters of pure luck, this would be to accept that the unlucky driver's and the successful would-be murderer's blameworthiness can depend entirely on luck.

Those who endorse the concessive view, on the other hand, could not accept that the lucky driver and the unsuccessful would-be murderer were not blameworthy at all, even though each would have been blameworthy had his actions resulted in a person's death. For concessivists deny that blameworthiness can depend *entirely* on luck. Assuming that each would indeed have been blameworthy had his actions resulted in a person's death,

[3] At least, had their actions resulted in a person's death in the way that Nagel is imagining that they might have. Their actions might also have resulted via some more convoluted and less foreseeable causal process in someone's death, and for this death we might want to deny that they were blameworthy. For example: a bystander might suffer a fatal heart attack upon hearing the screeching of the driver's tyres as he swerves onto the sidewalk, even though the driver hits no one. But I disregard this complication here and in what follows.

the concessivist would have to be committed to the view that their failure to kill did not absolve them of blameworthiness altogether. They would be blameworthy anyway in light of their respective recklessness and murderous intent. But each would have been *more* blameworthy had someone been killed as a result of his actions—even though the difference would have been a matter of pure luck.

Both extremists and concessivists conceive of intentional or careless risk imposition, then, as a kind of gamble with one's blameworthiness (cf. Jensen 1984: 326–7; Otsuka 2009: 375–7). But whereas extremists think that one can get away with the gamble scot-free, concessivists are more accommodating of the natural thought that even if one is lucky, and the consequences of one's intentional or careless risk imposition are not bad, one may nevertheless be blameworthy.

Both extremists and concessivists are opposed to what we can call

> *The sceptical view*: that an act's blameworthiness cannot depend on luck. (cf. Nagel 1979: 24)

The idea that underpins the sceptical view is recognisably Kantian, although Kant himself may not have accepted it.[4] The idea is that the only appropriate object of moral assessment is the quality of one's will, so that the quality of one's will alone determines the degree to which one is blameworthy for wrongdoing and the degree to which one should feel remorse.

[4] As Otsuka (2009: 385) notes, "Kant himself might have no quarrel with the claim that the shooter who kills has done a more serious wrong, is worthy of more severe punishment, and is perhaps even more morally blameworthy than the shooter who does not kill."

With respect to Nagel's examples, sceptics would argue that the intuitive judgments that Nagel describes should not be taken at face value. The lucky driver, who hits no one, is *just as blameworthy* as he would have been if he had hit someone, and should feel *just as remorseful*. Meanwhile, the unsuccessful would-be murderer is also *just as blameworthy* as he would have been had he succeeded in his murderous attempt, and should feel *just as remorseful*.

It is precisely because there is something intuitively appealing about the sceptical view that the reactions that Nagel describes give rise to a puzzle. As Nagel says,

> it is intuitively plausible that people cannot be morally assessed for what is not their fault, or for what is due to factors beyond their control … Without being able to explain exactly why, we feel that the appropriateness of moral assessment is easily undermined by the discovery that the act or attribute, no matter how good or bad, is not under the person's control … So a clear absence of control, produced by involuntary movement, physical force, or ignorance of the circumstances, excuses what is done from moral judgment. But what we do depends in many more ways than these on what is not under our control—what is not produced by a good or a bad will, in Kant's phrase. (1979: 25)

The sceptic's thought, then, is that when we have no control over some outcome, we are not to blame for producing it. (Nagel [1979: 26] calls this the 'condition of control'.) Yet the moral-luck supporting judgments that

Nagel describes are, for all that, themselves intuitive, familiar, and powerful. So, it seems that something has to give. We must repudiate either the sceptical intuition or our moral-luck supporting intuitions. We cannot keep both.

3. *A sceptical strategy*

One way to resolve the puzzle of moral luck is to provide compelling arguments for the appropriateness of the intuitions that tell in favour of her preferred resolution. Sceptics may try to give a deeper explanation, for example, for the indispensability of the condition of control. In light of the seemingly moral luck-supporting intuitions about Nagel's examples, the condition's intuitive plausibility isn't sufficient on its own to vindicate the sceptical view. But combined with a compelling deeper explanation, such as the one that we might find in the autonomy-based ethical system that Kant describes in the *Groundwork*, it might be. Similarly, believers in moral luck may try to describe a moral theory that gives a place to luck in determining blameworthiness that is so compelling that we are willing to discount the weight that we give to the condition of control. Martha Nussbaum (1986), for example, argues for an Aristotelian conception of morality and the good life according to which to be immune to the effects of luck in the assessment of one's character would be to be cut off from much of what is necessary for excellence, including moral excellence.

Another strategy for resolving the puzzle is to provide a debunking explanation of the intuitions that seem to support the position to which one is opposed. So, for example, sceptics may try to show that the seemingly moral luck-supporting intuitions that Nagel describes are not really to be

analysed in the way that I analysed them above. And believers in moral luck may try to show that the condition of control fails to capture exactly what it is that is intuitive in the idea that "people cannot be morally assessed for what is not their fault".

The most successful defence of one's favoured view will, most likely, offer a combination of these two strategies. What I have to say, however, is about the second strategy alone, and in particular on one version of that strategy that a number of sceptics have offered in support of their favoured resolution of the puzzle. I want to argue that it is not viable.

The strategy depends on a distinctive analysis of seemingly moral luck-supporting intuitions. Although the analysis in question is only one of a number that might be consistent with the sceptical view, I'll call it 'the sceptical analysis' for the sake of simplicity. The sceptical analysis has been pursued by a number of philosophers, most notably Brian Rosebury (1995) and, more recently, David Enoch and Andrei Marmor (2006) and Derek Parfit (2011).[5]

The idea, then, is to show that the seemingly moral luck-supporting intuitions that Nagel describes are not really to be analysed in the way that I analysed them above. They are to be analysed, the sceptic argues, as follows instead. In each case it is appropriate that in the event that the actions of the agent in question result in a death, the agent should undergo *worse experiences* than those that he should undergo in the event that his actions do not result in a death. But this isn't because the agent ought to feel greater *remorse* in the event that his actions result in a death. Rather, it is because there are good reasons for a person to suffer bad experiences *other than remorse* in the event that her blameworthy actions have bad consequences.

[5] Others suggesting this strategy include Richards (1986: 208); Thomson (1989: 213); Statman (1993: 16–7).

And these reasons do not apply in the event that her blameworthy actions do not have bad consequences—*even though she is just as blameworthy for her actions*. Thus, intuitions that seemed at first to support the possibility of moral luck turn out, on closer inspection, to be consistent with the sceptical view.

What is the nature of those bad experiences other than remorse? The sceptical analysis says different things about the careless driver and would-be murderer cases. So, let's begin with the former. Here, the analysis draws on a distinction between remorse, which is, on the framework that we are using, the appropriate response to only one's own blameworthiness, and the feeling or state that Bernard Williams identifies as 'agent-regret'.

As Williams describes it (see 1976: 123–6), agent-regret constitutively involves the thought 'how much better if it had been otherwise' that is characteristic of all forms of regret, and we can assume that, like all forms of regret, it pains the person who experiences it. But it is distinguished from regret in general by two things: its subject-matter, viz., the regretful person's own past actions; and the way in which it gives rise, in recognition of the regretful person's special relation to the victims of harms that her actions have brought about, to "a desire to make reparation" (Williams 1976: 125)—a desire that it may be appropriate to feel even if one is not to be blamed for the actions that brought about the harm, and which persists or would persist even were all costs associated with the harms already covered by insurance, for example.

The most important point about agent-regret for Williams's own purposes is not that agent-regret, as one's most basic feeling about a situation, may be appropriate even in cases where one's harmful actions were involuntary or innocent, but that it may *not* be inappropriate to *fail* to experience it, as one's most basic feeling about a situation, even in cases

where one's harm-producing actions were *wholly* voluntary under that description (and not justified by appeal to countervailing moral considerations or shielded from criticism by any moral permission). For Williams aims to bring out the ways in which whether one has cause to feel it may depend upon whether life-defining projects that one undertook and that could be expected to cause others harm succeeded, even though whether they would succeed was something that one could not control or be sure of. Thus, he thinks, someone who abandons his family to make a life as a painter may appropriately fail to feel agent-regret about the harm that he causes them if he succeeds, and yet only inappropriately fail to feel it if he fails. And this, Williams supposes, shows that a Kantian account of moral value, which has it both supreme and immune to luck, is mistaken (see Williams 1976: 116, 133–4).

In light of this, it may be a bit surprising to find sceptics about moral luck drawing on the notion of agent-regret in their sceptical analyses of intuitive judgments of moral luck. They do so by exploiting an ambiguity in Williams's discussion. That discussion doesn't make it altogether clear whether we should regard agent-regret among an agent's most basic feelings about a situation in which she knowingly and voluntarily brought harms about to be, when appropriate, a reflection of accurate *moral* assessment of her conduct or not. If it is, then the conclusion that it may not be inappropriate to *fail* to feel it even about such situations suggests that moral value is not immune to luck, because how one's conduct is to be morally assessed can depend on how things turn out that are a matter of luck. If it isn't, then the conclusion that it may not be inappropriate to fail to feel it even about such situations suggests that moral value is not supreme, provided that we make the further assumption—as Williams implicitly does (see especially 1976: 134)—that the rationally acceptable absence of agent-

regret as one's basic feeling about a situation implies the rational acceptability of failing to wish that one had acted otherwise.

Sceptics draw on the notion of agent-regret and yet avoid Williams's conclusions by repudiating any connection between agent-regret and moral assessment. Thus, the rational appropriateness of agent-regret is treated as independent of accurate moral assessment, so that there can be variation in the rational appropriateness of feeling it without variation in moral assessment, and vice-versa. And there is no need for sceptics to take a stand on Williams's implicit assumption about the implications of agent-regret and its absence for wishing that one had acted otherwise—although those sceptics who are influenced by Kantian thinking will be disposed to repudiate that too—since they are interested primarily in the relations between moral value and luck, not those between moral value and other forms of value.

Armed, then, with the distinction between agent-regret thus understood and remorse—which is, of course, not independent of moral assessment—the sceptic can press ahead with her analysis of careless driver case. As she points out, the unlucky driver, who kills the child, has reason to feel agent-regret. But he would not have had reason to feel agent-regret had he been lucky, and killed no one. So, all other things equal, the unlucky driver who kills the child has reason to *feel worse* than he would have had he been lucky. But both the lucky driver and the unlucky driver were careless to the same degree, and that carelessness, the sceptic claims, makes the same degree of *remorse* appropriate for both of them.

According to the sceptical analysis, then, the right way to understand our intuition about the careless driver case is as supporting the judgment that the unlucky driver ought to feel worse. But it is a mistake to understand it, with Nagel, as supporting the judgment that the unlucky driver ought to

feel more remorse because he is more blameworthy. That would be to conflate agent-regret with remorse. As Parfit writes, the lucky and unlucky drivers 'are both blameworthy, and should feel remorse', since both were careless. But '[it] is not clear that, as well as feeling…great agent-regret, the man who killed this child ought also to feel greater remorse.' (2011: 461; cf. Thomson 1989: 213; Rosebury 1995: 513–14)

Let's now turn to the would-be murderer case. Here, as we've seen, believers in moral luck may have the intuition that 'murder can be plausibly regarded as more blameworthy than attempted murder' (Parfit 2011: 157)—even where the difference is a matter of luck. Once again, the sceptical analysis holds that our intuition about this case is best understood as supporting the judgment that successful murderers ought to suffer worse experiences than unsuccessful would-be murderers. But, as before, that needn't be because of the appropriateness of greater *remorse*, or of subjecting him to greater *reproach*. Instead, we can appeal to the appropriateness of two bad experiences other than remorse and reproach. These are the agent-regret that a successful murderer ought to feel and the more severe punishment that she ought to receive.

The latter can be justified without appeal to greater blameworthiness. It may be appropriate, for example, because of the deterrent value of more severe punishment. If successful murders are punished no more severely than attempted murders, then 'when people have unsuccessfully attempted murder, they would have more reason to go on trying, until they succeed [since this] would not increase these people's punishment, if they were caught and convicted, and would often make it less likely that they would be caught.' (Parfit 2011: 462; cf. Rosebury 1995: 523) So, as Parfit concludes, it may 'sufficiently describe the moral difference' between the successful murderer and the unsuccessful would-be murderer to note that the former

ought, unlike the latter, to feel agent-regret and to be more severely punished. This is what lends Nagel's interpretation of our intuitions about the would-be murderer case its plausibility. But that interpretation conflates remorse and reproach with agent-regret and non-desert-based punishment. As with the careless drivers, both would-be murderers are blameworthy, and ought to feel greater remorse. But the successful murderer ought also to feel agent-regret and to suffer greater punishment.

4. *A challenge*

The sceptical analysis of the would-be murderer case may be right, although, in the absence of any other reason to doubt a belief in moral luck, it's not obvious that it's a *better* interpretation of our intuitions than Nagel's. But in any case, to vindicate the sceptical view, the sceptical analysis has to be a plausible analysis of *both* types case, since each constitutes independent support for the possibility of moral luck. I shall argue that it does not offer a plausible analysis of the careless driver case.[6]

The seemingly moral luck-supporting judgment about the careless driver case was, then, that the lucky driver is *less blameworthy* than the unlucky

[6] My strategy is not the one adopted by Domsky (2004), who also argues that the sceptical analysis is not a plausible one. His argument is that the analysis is unfaithful to one of the intuitions that generate the puzzle of moral luck in the first place, since the content of the relevant intuition is, precisely, that the unlucky driver ought to feel more *remorse*, and not merely *worse*, than the lucky driver. Domsky further suggests that scrutiny of the writings of those offering the sceptical analysis shows that they themselves accept this point. However, it is precisely the sceptic's point that how things may *seem* here is not how they actually *are*, and, as Statman (2005: 430–2) points out, the evidence that Domsky adduces in support of his further suggestion is not compelling.

driver, and should feel *less remorse.* As we saw, the sceptic's debunking explanation of this judgment charges its maker with conflating agent-regret and remorse in her analysis of our intuitions about the case. Both drivers, the sceptic says, should feel remorse, and it's not clear that one should feel more than the other.

But consider the following three versions of the careless driver case:

> *Bad Accident*: Distracted because I am talking on my mobile phone as I drive, I drift onto the wrong side of the road as you join it carefully from a side road. I hit and kill you.
>
> *Near Miss*: Distracted because I am talking on my mobile phone as I drive, I drift onto the wrong side of the road as you join it carefully from a side road. At the last moment, I realize what's happening and swerve out of the way. No one is hurt.
>
> *Empty Road*: Distracted because I am talking on my mobile phone as I drive, I drift onto the wrong side of the road. Fortunately, there are no other people around, and after a while I realize what's happening and return to the right side. No one is hurt.

Now, what the sceptic says about the careless driver thought experiment may seem plausible insofar as we imagine only a comparison between *Bad Accident* and *Near Miss*. Perhaps it will be agreed that I should feel the same degree of remorse, and am blameworthy to the same degree, in both of these cases. After all, as we might say, it's no thanks to the quality of my will or

intentions in *Near Miss* that you weren't killed. In that respect the comparison between these two careless driver cases is like that between the two murder cases. Intuitions that appear to support the views of believers in moral luck may turn out not to support them once we distinguish between the appropriateness of remorse and the appropriateness of agent-regret.

Insofar as we take the careless driver thought experiment to compare *Near Miss* with *Empty Road*, however, what the sceptic says about it is less plausible. I have no cause for agent-regret in either of these cases, since neither involves terrible consequences that arise as a result of my actions. (We can imagine that in *Near Miss* you are wholly unaware that I nearly killed you.) But it is nevertheless very plausible to think that I *wrong* you in *Near Miss*, and that this constitutes a morally significant difference between the two cases that justifies greater remorse (cf. Kumar 2003: 103). In *Empty Road*, even though I am careless to the same degree, I am lucky not only in that I *hit* nobody, but also in that I *wrong* nobody. (We can assume that I have no more reason to think that the road would be empty in *Empty Road* than I do in *Near Miss*.)

What does it add to say that I wronged you in *Near miss*? Even though you never know how close I came to hitting you, even though you suffer no harm at all as a result of my careless driving, your presence alters the moral significance of my action. It transforms it into a wronging, just as my taking an inaccurate pot-shot at what I unthinkingly took to be a scarecrow would be transformed from a mere wrongdoing (because I didn't bother to check before I took my pot-shot) into a wronging were the scarecrow in fact a farmer. One implication of this is that a third party who witnessed *Near Miss* would have good reason to feel indignation to a greater extent than would a third party who witnessed *Empty Road*. Another is that if you came

to learn of the events of *Near Miss*, you would have good reason to feel resentment to a greater extent than would anyone who came to learn of the events of *Empty Road*.

Of course, that I wrong nobody in *Empty Road* doesn't show that I do no wrong: a third party might still have *some* reason for indignation. But it's plausible to think that the appropriate response – from others and from the wrongdoer herself – to wrongdoing that also wrongs someone differs from the appropriate response to wrongdoing that wrongs no one. The same resolutions on my part concerning my future driving might be appropriate in both *Empty Road* and in *Near Miss*, but there seems to be much less reason for *remorse* in the former than in the latter – even if that is not to say that there is none.

This point is distinct from the one made by Nagel in his original article on moral luck and by many others since, viz., that unlucky consequences can justify greater remorse in cases of negligence. That's simply the putatively moral luck-supporting intuitive judgment. The contrast between *Bad Accident* and either of the other two cases may seem to believers in moral luck to illustrate that point, indeed, but sceptics will already be familiar with the contrast and ready to explain away the intuition in question as a conflation of remorse and agent-regret, as we saw. The contrast between *Near Miss* and *Empty Road* is different because *there are no consequences that would constitute grounds for agent-regret in either case*. Yet the intuitive difference in the remorse that's appropriate in the respective cases remains.

If this analysis of *Near Miss* and *Empty Road* is accurate, then the sceptic's attempt to diagnose apparently moral luck-supporting intuitions about the careless driver thought experiment as based on a mistaken conflation of remorse and agent-regret turns out at best to be much less powerful than it might seem at first. Insofar as we imagine the thought experiment to

involve a comparison of something like *Bad Accident* with something like *Near Miss* or *Empty Road*, the sceptic's debunking analysis may seem plausible. But the possibility of a version of the thought experiment that compares *Near Miss* with *Empty Road* undermines the analysis's ability to handle risky behaviour cases in general. When there is no cause for agent-regret in either scenario, apparently moral luck-supporting intuitions cannot be so easily explained away. Even if the sceptic continues to endorse her analysis of the comparison between *Bad Accident* and *Near Miss* or *Empty Road*, then, that analysis no longer appears sufficient to undermine belief in resultant moral luck *in general*.

5. *Objections*

One objection to the foregoing argument concerns my claim that there are no consequences that would constitute grounds for agent-regret in either *Near Miss* or *Empty Road*. According to the description and analysis of the cases that I offered above, a possible consequence of my careless driving in virtue of which I turn out to be more blameworthy is that I *wrong* you. Could it not be that when that consequence eventuates, as it does in *Near Miss*, it constitutes grounds for agent-regret after all? If it could, then my analysis of the comparison between *Near Miss* and *Empty Road* is mistaken: the cases differ in that the former and not the latter is a case in which my behaviour has consequences that would constitute grounds for agent-regret, and so the intuition that I have reason to feel more remorse in the former than in the latter may, after all, be explained away as the sceptic suggests.

It is implausible, however, to include the fact of having wronged someone among the possible grounds for the reaction of agent-regret, as the

sceptic conceives it. A wronging, I take it, is a species of wrongdoing, rather than something over and above wrongdoing that may sometimes accompany it. Indeed, it's misleading to conceive of wronging as a *causal* consequence of certain wrongdoings. Instead, the right way to think about wrongings is as follows: a logical consequence of the conjunction of certain wrongdoings and their contexts and causal consequences is that those wrongdoings are wrongings.

The appropriate reaction to any species of wrongdoing in itself is remorse; the degree of remorse that is appropriate is a function of the seriousness of the wrongdoing; and the seriousness of the wrongdoing is, I claim, affected by whether or not it turns out to be a wronging. To think that agent-regret would be appropriate in the case of a wronging such as that of *Near Miss* would be to suppose that we could disentangle the wrongdoing and the wronging in order to make each the subject of a different reaction, for there is, as I have argued, no other possible object for the reaction of agent-regret. Since this is a mistake, this objection fails.

A second objection takes issue with my characterisation of *Empty Road* as a case in which I wrong no one. The mere fact that I don't come close to hitting anyone, you might think, doesn't suffice to show that I don't wrong anyone. Among those whom I wrong in driving so carelessly, for example, are the people who might easily have been on the road, such as those who normally use the road at around that time of day. I wrong these drivers even though they weren't, as it happens, using the road at that time on this particular day.

If this is right – if I wrong someone in *Empty Road* just as I do in *Near Miss* and *Bad Accident* – then my analysis of the comparison of *Near Miss* and *Empty Road* won't work. It won't, that is, support the conclusion that I am

more blameworthy in *Near Miss* on account of the fact that I wrong someone in that case whereas I don't wrong anyone in *Empty Road*.

However, it isn't plausible to say that in *Empty Road* I wrong the people who might easily have been on the road. One appropriate reaction to one's having wronged someone is to apologise to her. But it would be very odd for me to go around apologising to all the people who might easily have been on the road at the time of my carelessness. After all, if I wrong those people in *Empty Road*, presumably I wrong them also in *Bad Accident*, since in this cases it's only a matter of luck that I didn't hit (or nearly hit) them instead or as well as you, the victim of my carelessness. But certainly no one would think that I ought to go around apologising to them as well as making any appropriate reparations to your family, for example.

A third objection puts pressure on the claim that in *Near Miss* I wrong you. Suppose that it's true that in *Empty Road* I wrong no one. I do wrong, of course, in driving carelessly without particular reason to think that my doing so subjects no one else to risk of injury. But, as I just argued, there is no one whom I wrong: all the people who might easily have been on the road, such as those who normally use the road at around that time of day, happened not to be.

The challenge, now, is to explain how your position in *Near Miss* differs from the position of those in *Empty Road* who might easily have been on the road but were not in actual fact. After all, neither you nor they were harmed by my carelessness; indeed, as I suggested earlier, we can imagine that no one beside me even knew of it. The only difference seems to be that you were physically nearer to me, so that my carelessness would have caused you harm had I not noticed what was happening and swerved. But surely that's not morally significant enough to make the difference between wronging and not wronging?

Such a difference *is* morally significant enough. But that's not because you cross, in *Near* Miss, some independently morally significant threshold of physical proximity that's not crossed by all the people who might easily have been on the road but were not in actual fact. Rather, it's because in *Near Miss*, you are actually exposed to the risk of harm by my carelessness: your physical proximity places you within the scope of those who are so exposed. Other people, people who might easily have been on the road but are not in actual fact, are not actually exposed to the risk of harm by my carelessness. They are exposed to that risk only in nearby possible worlds in which they *are* on the road.

It's a morally significant fact that I expose you to the risk of harm in *Near Miss*, and it constitutes a morally significant difference between you and others that you are so exposed and they are not. That's why I wrong you but not them. Thus, this third objection fails.

The final objection that I'll consider raises the possibility that something like the sceptic's debunking analysis is available after all to explain away the seemingly moral luck-supporting intuitions that I've described regarding the comparison between *Near Miss* and *Empty Road*.[7]

The sceptic argues that in *Bad Accident* I ought to feel worse than I ought to feel in *Near Miss* because agent-regret would be appropriate in the former but not in the latter. As we've seen, agent-regret is a painful appreciation of one's relation to the victims of harms that one's actions have produced, which manifests itself in a desire to make reparation. The sceptic further specifies that the rational appropriateness of agent-regret varies independently of blameworthiness.

---

[7] Thanks to Michael Otsuka for this suggestion.

As I've stressed, agent-regret so understood wouldn't be appropriate in either *Near Miss* or *Empty Road*, since my actions produce no harms in either. Yet the desire to make reparations may nevertheless be appropriate in *Near Miss* but not in *Empty Road*, and this, the sceptic might argue, signals the presence of a painful feeling *analogous* to agent-regret—call it agent-regret*—even if not identical with it, which varies independently of blameworthiness and which can, as a result, supply the basis of a debunking analysis of seemingly moral luck-supporting intuitions about the comparison between the two cases. I ought to feel worse in *Near Miss* than in *Empty Road*, the sceptic can say, but that's not because greater *remorse* would be appropriate. It's because agent-regret* would be appropriate in *Near Miss*, but not in *Empty Road*. The two cases do not differ in respect of my blameworthiness, but they differ in that in *Near Miss* I wrongfully exposed someone to the risk of terrible consequences through my agency, and in *Empty Road* I did not. Hence, the intuitive—but, as we now see, mistaken—impression of a difference in blameworthiness between the two cases.

However, this strategy is less plausible than the sceptic's debunking analysis of the intuitions about the comparison between *Bad Accident* and *Near Miss* or *Empty Road*. The key thing for the sceptic is to ensure that agent-regret* can be distinguished from remorse. The sceptical strategy as applied to that comparison, as we saw earlier, depends on a clear separation between moral assessment and agent-regret; where Williams leaves it unclear whether the rational appropriateness of agent-regret is bound up—at least in the realm of voluntary actions—with the rational appropriateness of moral disapproval, the sceptic must stipulate that it is not. (Otherwise, one can be appropriately subjected to moral disapproval even for things that were beyond one's control.) But if it is not, then it must have as its object—it

must be construed as a response to—something more than the mere fact of having wronged others. For the rational appropriateness of a painful appreciation of one's relation to the victims of *wrongs* that one has done surely cannot vary independently of the rational appropriateness of moral disapproval. Such an appreciation is surely indistinguishable from *remorse*.[8] So, it's important for the sceptic to understand agent-regret as a response to *harms*, rather than to *wrongs*. That makes sense, too, of Williams's point that agent-regret can be appropriate even when one was not at fault for the production of the harms.

Now, if our intuitions are analysed as reflecting the appropriateness of greater remorse in *Near Miss* than in *Empty Road*, the sceptical strategy fails, since its very purpose is to undermine the impression of that appropriateness. So, what the sceptic means by agent-regret* had better be something more than a response to wrongs that one has done, which manifests itself in a desire to make reparation, just as what the sceptic meant by agent-regret proper had to be something other than such response. But if there are no *harms* around for agent-regret* to be a response to, as there aren't in *Near Miss*, then what else *could* it be a response to except the very fact of wronging itself?

I can think of nothing that could stand in the place of the *actual* harms that agent-regret proper is a response to. So, the possibility that the objection under consideration raises—the possibility of running an analogous debunking analysis of the appearance of a commitment to moral luck in our judgments about *Near Miss* and *Empty Road*—looks remote. At the very least, we must be given some indication of what agent-regret* is a painful consciousness *of*, if not the fact of having caused *harm* to someone

[8] Compare Williams's comments (1976: 126) about the way in which remorse is often understood, viz., as 'what we have called "agent-regret", but under the restriction that it applies only to the voluntary.'

(since there is no such harm in either *Near Miss* or *Empty Road*) or the fact of having *wronged* someone (since consciousness of that fact won't do the job that agent-regret* is being invoked to do). Since the possibility that the objection raises looks remote, I conclude that the objection itself is at best under-motivated.

6. *Conclusion*

My argument against the sceptic's debunking explanation isn't, of course, evidence for the possibility of moral luck. As we saw earlier, a full argument for that possibility would need to offer reasons for a restricted interpretation of the condition of control, at least, as well as compelling arguments in favour of moral luck-supporting intuitions. But, since the sceptic's debunking explanation is an attempt to undermine the appearance that those intuitions are in fact moral luck-supporting, my argument is constitutes a defence of the possibility of moral luck even if it's not evidence for that possibility.

References

Domsky, D. 2004. There Is No Door: Finally Solving the Problem of Moral Luck. *The Journal of Philosophy* 101: 445–464.

Jensen, H. 1984. Morality and Luck. *Philosophy* 59: 323–30.

Kumar, R. 2003. Who Can Be Wronged? *Philosophy & Public Affairs* 31: 99–118.

Nagel, T. 1976. Moral Luck. *Proceedings of the Aristotelian Society, Supplementary Volumes* 50: 137–51.

Otsuka, M. 2009. Moral Luck: Optional, Not Brute. *Philosophical Perspectives* 23: 373–387.

Parfit, D. 2011. *On What Matters*. Oxford: Oxford University Press.

Richards, N. 1986. Luck and Desert. *Mind* 65: 198–209.

Rosebury, B. 1995. Moral Responsibility and 'Moral Luck'. *The Philosophical Review* 104: 499–524.

Statman, D. 1993. Introduction. In *Moral Luck*, ed. D. Statman, 1–34. New York: State University of New York Press.

Statman, D. 2005. Doors, Keys, and Moral Luck: A Reply to Domsky. *The Journal of Philosophy* 102: 422–436.

Strawson, P.F. 2003. Freedom and Resentment. In *Free Will*, ed. G. Watson, 72–93. Oxford: Oxford University Press.

Thomson, J. J. 1989. Morality and Bad Luck. *Metaphilosophy* 20: 203–21.

Williams, B.A.O. 1976. Moral Luck. *Proceedings of the Aristotelian Society, Supplementary Volumes* 50: 115–35.

Zimmerman, M. J. 1987. Luck and Moral Responsibility. *Ethics* 97: 374–386.