

MORAL AND EPISTEMIC LUCK

ABSTRACT. It is maintained that the arguments put forward by Bernard Williams and Thomas Nagel in their widely influential exchange on the problem of moral luck are marred by a failure to (i) present a coherent understanding of what is involved in the notion of luck, and (ii) adequately distinguish between the problem of moral luck and the analogue problem of epistemic luck, especially that version of the problem that is traditionally presented by the epistemological sceptic. It is further claimed that once one offers a more developed notion of luck and disambiguates the problem of moral luck from the problem of epistemic luck (especially in its sceptical guise), neither of these papers is able to offer unambiguous grounds for thinking that there is a problem of moral luck. Indeed, it is shown that insofar as these papers succeed in making a *prima facie* case for the existence of epistemic luck, it is only the familiar sceptical variant of this problem that they identify.

0. INTRODUCTION

A generation has now gone by since the original exchange between Bernard Williams (1976) and Thomas Nagel (1976) that prompted the contemporary discussion about moral luck, and now that we have such temporal distance between the original stimulus for the contemporary debate and the debate itself it is worthwhile re-examining this exchange in a new light.¹ For what seems to be widely acknowledged about these papers is that at the very least they draw our attention to a genuine tension in our moral concepts. It is this claim that I want to call into question here, on two grounds. The first is that the failure on the part of either of these philosophers to put forward even a partial account of luck disguises the fact that by the lights of any plausible account of this notion the examples of moral luck that they offer are ambiguous, to say the least. The second is that both of these papers illicitly appeal to the analogue problem of epistemic luck, and that with this sister problem disambiguated from the problem of moral luck the *prima facie* force of the examples of moral luck that are presented is further undermined.

Note that my claim is not that there is (or could be) no coherent account of the problem of moral luck, but rather the more modest contention that the arguments offered by the two papers that form the *locus classici* as regards the contemporary debate on this subject matter are inadequate to the purpose. Clearly, however, the modest contention suggests the grander claim even if it does not entail it.

1. MORAL LUCK AND LUCK *SIMPLICITER*

Nagel describes the phenomenon of moral luck as being concerned with the following types of scenario where

[...] a significant aspect of what someone does depends on factors beyond his control [*and*] yet we continue to treat him in that respect as an object of moral judgement [...]. (Nagel 1979, 25)

So, for example, the consequences of my decision to reverse out of my driveway without checking first to see if anyone is approaching along the pavement are, to a certain extent at least, subject to luck. Nevertheless, I will be held to account for those consequences even if they are the result of luck. If, for example, it just so happens that an elderly gentleman is walking past my driveway as I reverse out of it, and I hit him and kill him, then (intuitively at least) I will be subject to a greater level of moral opprobrium than if he had not happened to be there and I had exited my driveway without incident. It thus seems that my moral responsibility extends even to ‘lucky’ factors that I have no control over, and this is contrary to a widespread intuition that (roughly) we are only morally responsible for what is in our control.

While there has been a tremendous amount of critical appraisal of the general notion of moral luck, what has been noticeably lacking in the literature is an assessment of the manner in which Williams and Nagel employ the concept of luck in their arguments.² Neither of them offer an explicit account of the notion, with Williams remarking (almost in passing) that he will “use the notion of ‘luck’ generously [*and*] undefinedly” (Williams 1981*b*, 22). Commentators have followed suit and thus been likewise inclined to treat the concept as an undefined primitive that we may legitimately take ourselves as having a clear grasp of. As we will see, however, the notion of luck that is at issue will have ramifications for the plausibility of the examples that Williams and Nagel present in favour of their claim that there is a problem of moral luck, so it is necessary to look a little deeper at what kind of notion of luck is being presupposed here, if any.

To begin with, note that in the quotation just cited all that we can discern is that Nagel thinks that there is a close connection between luck and lack of control. This point about luck being associated with lack of control is echoed in one of the few passages in the Williams paper where he mentions luck in isolation from moral luck. He writes that “what is not in the domain of the self is not in its control, and so is subject to luck” (Williams 1981*b*, 20).³ Clearly, however, this could only be (and is only intended to be) a very partial account of luck since all sorts of events—such as the celestial movements of the planets—are beyond our control, but are not thereby considered lucky as a result.

So how then should we understand luck, if it isn’t to be simply identified with events that one lacks control over? I think that we can get a rough approximation to what is involved in the notion of

luck via the following characterisation:

Luck

A lucky event is an event that occurs in the actual world but which does not occur in most of the nearest possible worlds to the actual world where the relevant initial conditions for that event are the same as in the actual world.

Consider, for example, how this understanding of luck will deal with paradigm cases of lucky events, such as a lottery win. Here we have an event that obtains in the actual world but which does not obtain in most of the nearest possible worlds to the actual world where, intuitively, the relevant initial conditions are the same. Although one might win a free and fair lottery with long odds in the actual world, in most near-by possible worlds where, for example, one continued to buy the lottery ticket that one did and the lottery remained free and fair, one would not be standing here now clutching the winning lottery ticket. Moreover, this account will also explain why there is the close connection noted by Williams and Nagel between an event being lucky and that event being out of the agent's control. If one is unable to control whether or not a certain event obtains then this will mean that while this event might well obtain in the actual world, we would not expect it to obtain in near-by possible worlds where one's inability to control this outcome is kept constant. My winning of a boxing match against the world champion is out of my control, and thus lucky, on this account, because while I might as it happens win in the actual world (where my opponent has the misfortune to sprain his ankle as he enters the ring, thus forfeiting the fight), in most near-by possible worlds where the relevant initial conditions are the same (in particular, where our fighting abilities are kept constant), I will be thrashed to a pulp.

This is not to suggest, of course, that this is a fully-fledged characterisation of the notion since it is still somewhat vague, especially regarding what counts as the 'relevant' initial conditions. Moreover, it seems that we would need to supplement such a formulation with a further element so as to account for the fact that it is only events that are in some way significant to the agent concerned that are counted as lucky (or, indeed, as unlucky). As this characterisation of luck stands, all manner of 'odd' events—such a freak landslide on one of the moons of Jupiter—might fit the rubric for luck which no one would regard as being of any significance and thus as being either lucky or unlucky. Nevertheless, such a proposal is manifestly on the right tracks and a clear improvement on the cursory remarks offered by Williams and Nagel.⁴ With this account of luck in mind, we will consider two of the key examples of moral luck that Nagel offers.

2. NAGEL ON MORAL LUCK

Perhaps the most famous example that Nagel offers in favour of moral luck is that of the drunk driver. Nagel writes that

[...] there is a morally significant difference between reckless driving and manslaughter. But whether a reckless driver hits a pedestrian depends on the presence of the pedestrian at the point where he recklessly passes a red light. (Nagel 1979, 25)

And since the presence of the pedestrian at that particular point and at that particular time is a matter of luck, so luck alone can make a morally significant difference (the difference between (mere) reckless driving and manslaughter). We are thus asked to imagine two drunk drivers who are counterparts in every respect except that the former has the bad luck to hit a pedestrian while the latter has the good luck not to, and where these differences in consequences are morally significant.⁵

Of course, simply noting the difference between the *crimes* of reckless driving and manslaughter is not enough to establish the point that Nagel wishes to argue for. It is not in itself contentious to suppose that two agents could be equally morally at fault (and otherwise identical) and yet the one agent be guilty of a lesser criminal offence (and so subject to a less severe punishment). There are a number of reasons why this might be so. For one thing, the role of punishment is standardly understood such that it should do more than merely reflect moral disapproval (indeed, some might argue that punishment should not reflect moral disapproval at all). For similar reasons, it is even less obvious that the extent of the punishment should match the extent of the moral opprobrium in each case. Moreover, punishments can vary in line with there being a victim to the crime in question (and thus vary in response to the extent of the suffering of the victim). Whether rightly or wrongly, we might wish to punish one criminal more severely than another for committing the same crime for the sole reason that the first criminal's act, while otherwise identical, resulted in more suffering and we want our punishments to (somehow) represent this differing extent of suffering. Since there is this logical gap between moral opprobrium and punishment, the onus is on Nagel to do more than merely show that there is a difference in terms of the crimes that we attribute to the agents in question (and thus the punishments we inflict) if he is to show that there is a genuine moral difference in these cases.

One response to the putative examples of moral luck that Nagel and others have proposed has thus been to put pressure on this potential gap between moral opprobrium and legal punishment (in the form of the crimes attributed to the agent, and the punishments incurred by the agent as a result). This style of counterargument essentially involves taking the examples offered at face value and trying to show how the intuitions that they give rise to can be reinterpreted so that our initial 'intuition' of there being a moral difference can be explained away.⁶ In contrast, I contend that the focus of our

counterargument should, as it were, be one stage back by not taking these examples of moral luck at face value in the first place. The claim is that such examples are, on closer inspection, controversial because they trade on diverse—and, in this case, incompatible—claims about luck. As we will see, once the role of luck in these examples is made clear we will have a way of responding to the putative phenomenon of moral luck that both supplements and strengthens the more conventional critique offered by those who merely note the gap between moral opprobrium and legal punishment.

There are many possible ways of filling-out the details of Nagel's drunk driver example and, depending on what detail we add, the example may not end up involving luck at all (at least regarding whether or not the agent hits the pedestrian). Indeed, we can imagine two extreme cases, one where there is a lot of luck involved and one where there is hardly any (and perhaps none at all), with most other cases lying on a continuum between these two extremes. At the 'non-lucky' end of the continuum will be those cases where the agent in question regularly takes risks of this sort and where he is taking a great risk in driving under the influence like this, perhaps because he is about to drive down a crowded street. If this is how we are to understand the example, however, then it would be odd to say that the driver was unlucky to hit a pedestrian, since we would *expect* him to hit a pedestrian, not just in this world but in all near-by possible worlds where the relevant initial conditions were the same (where, for example, his reckless character remained unchanged). At the other end of the continuum is the genuinely unlucky drunk driver who isn't taking much of a risk (it is a rural area, say, with very few pedestrians), and who rarely takes risks of this sort (his reckless behaviour is, we might say, out of character). Here we do have a case of bad luck since although the driver hits a pedestrian in the actual world, we would not expect him to hit a pedestrian in most near-by possible worlds where the relevant initial conditions are the same. In these worlds, we would expect him to either come to his senses and not drive, or else drive and not hit a pedestrian.

So we can either understand the drunk driver example in such a way that he is clearly unlucky to hit a pedestrian, or we can understand it in such a way that the driver's actions having this consequence isn't a matter of luck at all. Clearly, there is a moral difference between these two drivers even if the consequences of their actions turn out to be the same. For example, if they both end up hitting a pedestrian then I think that we would subject the unlucky driver to a lesser degree of moral censure than his non-lucky counterpart. Part of the reason for this is that, unlike the non-lucky driver, the unlucky driver acted out of character. Moreover, he acted in such a way that was far less reckless than the actions of the non-lucky driver.

Of course, this kind of moral contrast between the unlucky and non-lucky drunk driver is of no use to Nagel because the surrounding facts of the situation are substantively different in each case. Accordingly, he will not be able to employ this contrast in order to motivate the desired conclusion that

luck alone is affecting the kind of moral evaluation that is being offered. It thus appears that if Nagel is to get the unambiguous example of moral luck that he wants, then he is going to have to consistently understand the drunk driver case along one of the two lines outlined above. In the one case (the driver who acted out of character) we would have an agent who was genuinely unlucky to hit a pedestrian, unlike his (non-lucky) counterpart. In the other, we would have an agent who was lucky to *not* hit a pedestrian, unlike his non-lucky counterpart. The issue thus concerns whether the two counterparts consistently understood in either of these ways should be subject to a different moral evaluation. We will take the cases in turn. First, the driver who is unlucky to hit a pedestrian, unlike his counterpart who doesn't.

Although there is obviously a difference in consequences here, and thus a difference in the crime that the driver will be charged with in each case (manslaughter in the former case, rather than just reckless driving), once we had reflected on the details of the scenario we would, I think, feel a certain sympathy for the unlucky driver that his lack of care should have had such tragic consequences. In particular, I think that we would feel sympathy for this driver precisely because of the fact that his actions were out of character and would not normally have led to anyone being harmed. That is, the fact that the agent is genuinely unlucky in having run over a pedestrian would temper our moral disapproval of him. That our response to the dire consequences brought about (in part) by luck should be one of sympathy itself suggests that we are willing to 'factor-out' the role of luck in our moral assessment of an agent's actions. While we might recognise the need for an appropriate legal censure which is greater than his counterpart, this need not reflect any conviction that the unlucky driver has committed a morally worse action than his counterpart.

Similar remarks apply to the other case, that of the drunk driver who was not acting out of character and who was taking a great risk by driving home drunk on roads where there are lots of pedestrians. Suppose that the agent in question has the good fortune to not hit a pedestrian, unlike his counterpart who, as we would expect, does. Would this lead to a different moral assessment of the two agents? Again, I think that we would answer in the negative, even though we would grant that the crimes committed by these agents (and thus the punishments that they should incur) would be different. If we knew that the agent who was lucky to not hit a pedestrian was acting in character in taking such an inordinate risk with other people's lives, then I think that we would regard him as behaving in a way that was no less subject to moral censure than his counterpart. Indeed, we could imagine, for example, this lucky agent being condemned in court by the judge for behaving in a way that was, morally, no different from someone who had actually run over a pedestrian (even while passing a different sentence to that which he would have passed to a counterpart driver who did run over a pedestrian).

So if Nagel wants to make use of our clear intuition that there is a moral difference in the drunk

driver example, then he needs to vary the details of the example so that we are making a comparison between the ‘out-of-character’ driver (or some analogue) and the ‘in-character’ driver (or some analogue). Crucially, however, this contrast involves more than just a difference in the luck involved in each case, since it also illicitly varies circumstantial features of the scenario at issue. Insofar as Nagel sticks to examples which do not vary circumstantial features of the scenario and simply varies the luck involved, however, as in examples that consistently stick to either the out-of-character or the in-character template, then he doesn’t get the clear moral intuitions that he is trying to motivate.

The same goes for the other examples that Nagel offers, though the details are different in each case. The example of the drunk driver is an instance of what Nagel calls “resultant luck”, which is “luck in the way one’s actions and projects turn out” (Nagel 1979, 27). This is contrasted with what he terms “circumstantial luck”, which is luck involved in the “kind of problems and situations one faces” (Nagel 1979, 27).⁷ The main example that he offers to illustrate this kind of luck is of the ‘unlucky’ Nazi and the ‘lucky’ German expatriate:

[*W*]hat we do is [...] limited by the opportunities and choices with which we are faced, and these are largely determined by factors beyond our control. Someone who was an officer in a concentration camp might have led a quiet and harmless life if the Nazis had never come to power in Germany. And someone who led a quiet and harmless life in Argentina might have become an officer in a concentration camp if he had not left Germany for business reasons in 1930. (Nagel 1979, 25)

We are clearly meant to suppose that it is the same agent in each of these cases, though faced with different situations. Accordingly, we are to imagine an agent who would have led a harmless life if the Nazis had never come to power, or if he had emigrated to Argentina for business reasons in 1930, but who in fact became an officer in a Nazi concentration camp.

In these, and other examples that Nagel offers, reflecting on the role of luck in the example does not just undermine the force of our initial intuitions in favour of moral luck, as it does in the drunk driver case, but actually completely undercuts them. For suppose we take seriously the idea that had our protagonist been ‘lucky’ to have avoided being present when the Nazis were in power, then he would have led a relatively blameless life. According to the understanding of luck offered here, this means that we have to suppose that there are a great number of near-by possible worlds in which this agent lives under the Nazi rule and so commits atrocities as a result. If this is to be taken completely at face value, then I think we would agree, on reflection, that there is no clear moral difference between the Nazi officer and his ‘peaceful’ Argentinean counterpart. Indeed, we often find out key moral truths about agents by getting a ‘glimpse’ of how they might have behaved had circumstances been different (we might observe, for example, their surprising degree of callousness when faced with an injured animal, or their cruel disregard for another’s feelings). Of course, part of the problem of evaluating such cases is that we have such a shaky epistemic access to the relevant counterfactual facts. Our

‘hunches’ about our ‘peaceful’ expatriate German neighbour in Buenos Aires are only that, and it is rare that we would have evidence of any definitive sort to justify such a damning verdict about someone who did not actually commit the crimes in question. Nevertheless, insofar as we are entitled to take the relevant facts as known, as Nagel implicitly asks us to do, then the putative moral difference disappears.⁸

In contrast, when the situation is described so that we are willing to suppose that there is a moral difference between the peaceful German expatriate and his Nazi counterpart, then luck no longer seems to be playing the desired role. For example, when it is stipulated that there are very few (if any) near-by possible worlds in which our naturalised Argentinean commits such atrocities, then our intuition that there is a moral difference between the peaceful expatriate German and his Nazi counterpart is re-established, but at the expense of this no longer being an example that illustrates moral luck. After all, since we have now stipulated that the possibility that the agent could have been a Nazi is remote, it follows that it is not a matter of luck that the agent leads the peaceful—as opposed to wicked—life that he does. Indeed, given the remoteness of the possibility, it is now a contentious issue whether we are talking about the same agent in each case (or, at least, relevantly similar agents). So either Nagel can get the relevant moral difference but in doing so loses the sense in which luck is involved, or else he retains the role of luck in the example but at the expense of completely undermining our intuition that there really is a moral difference between the two counterpart agents.

The reason why reflecting on the role of luck in this example has a more dramatic effect than it does in the case of the example of resultant luck considered above is that whereas the focus in the drunk driver case is on our moral assessment of the actions of the driver, in the expatriate case it is more on the agent himself. This makes a substantive difference because although we can (though with difficulty, as it turns out) conceive of scenarios in which two otherwise identical agents engage in the same act and yet, due to luck, that act seems to be subject to differing moral evaluations, it is far harder (if not impossible) to imagine two otherwise identical agents engaging in lifestyles that—again, due to luck—are so drastically divergent in their moral status. That is, to suppose that two agents are otherwise identical is to thereby suppose that while luck might influence their resultant behaviour in different ways, they are not, *qua* agents, subject to a different moral assessment as a result. So while it might be a matter of luck that an agent does not end up being an officer in a Nazi concentration camp, it will not be a matter of luck that the agent displays behaviour that is subject to moral censure (and indicative of a generally morally corrupt character) in other ways. Conversely, in cases such as the drunk driver example, although focusing upon the actions might seem to present some (albeit, as we have seen, inconclusive) grounds for thinking that luck can influence our moral assessment of an agent’s actions, it would still remain that we would hold the protagonist and his counterpart as equally

at moral fault *qua* agents. It is thus significant that Nagel motivates his case for the circumstantial luck that (putatively) affects our assessments of agents via the less contentious (though still problematic) case of resultant luck.⁹ Whereas the latter merely trades upon ambiguous claims about luck that, once disambiguated, weaken his argument, the former trades upon the supposed truth that there exist genuine cases of resultant moral luck in order to establish the even more contentious claim that there exist cases of circumstantial moral luck also.

3. WILLIAMS ON MORAL LUCK AND RATIONAL JUSTIFICATION

As we have seen, Nagel's examples fail to work because he doesn't keep the relevant details of the examples fixed. We cannot simply extend this critique of Nagel to Williams' own treatment of the issue, however, because Williams offers a very different account of what is involved in the putative phenomenon of moral luck. What Williams explicitly does (and which Nagel only does implicitly), is recognise that the problem of moral luck is in fact derivative on the problem of epistemic luck.¹⁰ After all, while, as we have just seen, this issue about lack of control over consequences does not seem to undermine the moral evaluations of actions, intuitively it *can* undermine the epistemic status of one's assessment of what the consequences of actions might be. Crucially, however, we have now shifted from the issue of how luck can undermine the moral status of actions to how it can undermine the epistemic status of beliefs. This point is brought out most clearly once one looks at Williams' paper on this topic, since for him the focus when it comes to moral luck is upon what he calls the "rational justification" of actions (Williams 1981*b*, 22), and thus his article is not obviously concerned with the specifically moral status of actions at all.

The primary example of moral luck that Williams offers in his article is that of the painter Gauguin who deserts his family in order to pursue his ambition to be a great artist (Williams 1981*b*, 22ff.). Williams points out that Gauguin cannot be sure in advance that his project will be successful, and yet the success or otherwise of the project will determine how we morally evaluate his decision to desert his family. Williams writes:

[W]hether he will succeed cannot, in the nature of the case, be foreseen. We are not dealing here with the removal of an external obstacle to something which, once that is removed, will fairly predictably go through. Gauguin, in our story, is putting a great deal on a possibility which has not unequivocally declared itself. I want to explore and uphold the claim that in such a situation the only thing that will justify the claim will be success itself. If he fails—and we shall come shortly to what, more precisely, failure may be—then he did the wrong thing, not just in the sense in which that platitudinously follows, but in the sense that having done the wrong thing in those circumstances he has no basis for the thought that he was justified in acting as he did. If he succeeds, he does have a basis for that thought. (Williams 1981*b*, 23)

So if Gauguin's action does result in him becoming a great painter then that action will be rationally justified (despite the cost of this action to his family and others), whereas if he does not succeed then his action will not be rationally justified.

Crucially, argues Williams, it can be a matter of luck that his action leads to success. Williams is clear, however, that not just any kind of failure will suffice to show that Gauguin's decision was unjustified. After all, as he points out,

[...] if Gauguin sustains some injury on the way to Tahiti which prevents his ever painting again, that certainly means that his decision (supposing it now to be irreversible) was for nothing, and indeed there is nothing in the outcome to set against the other people's loss. But that train of events does not provoke the thought in question, that after all he was wrong and unjustified. He does not, and never will, know whether he was wrong. What would prove him wrong in his project would not just be that *it* failed, but that *he* failed. (Williams 1981*b*, 25, *italics mine*).

What is required for failure is thus some genuine test of Gauguin's choice which shows that he made the wrong decision, rather than merely an external obstacle preventing that choice from even being tested in the first place, and Williams' point is that luck can intervene even here. Williams does not give an example of luck that is "intrinsic" rather than "extrinsic" to Gauguin's project, though he does offer a different example regarding Tolstoy's fictional heroine Anna Karenina:

Anna remains conscious in her life with Vronsky of the cost exacted from others, above all from her son. She might have lived with that consciousness, we may suppose, if things had gone better, and relative to her state of understanding when she left Karenin, they could have gone better. As it turns out, the social situation and her own state of mind are such that the relationship with Vronsky has to carry too much weight, and the more obvious that becomes, the more it has to carry; and I take that to be a truth not only about society but about her and Vronsky, a truth which, however inevitable Tolstoy makes it seem, could, relative to her earlier thoughts, have been otherwise. It is, in the present terms, a matter of intrinsic luck, and a failure in the heart of her project. But its locus is not by any means entirely in her, for it also lies in him. (Williams 1981*b*, 26-7)

This example of intrinsic luck suggests that intrinsic failure in the case of Gauguin's project would consist in Gauguin arriving in Tahiti and giving his painting the best shot that he can (unencumbered by extrinsic luck), but ending up with disappointing results nonetheless. Crucially, however, the failure at issue here is not meant to reflect a *mistake* on Gauguin's part regarding his original assessment of his abilities, since we are meant to suppose that, relative to his "state of understanding" when he made his decision, that decision was based on entirely epistemically justified grounds. Rather, the failure consists in something other than that, though intrinsic to the project nonetheless. Finding uncontroversial examples here is difficult (which is probably why Williams does not even try), but one possibility might be that what Gauguin eventually discovers in Tahiti, contrary to the information that he had to go on while back in Europe, is that his abilities were in fact dependent upon the stresses and hardship that life with his family brought him.

With the problem so construed, however, it is far from clear that there is a problem regarding moral luck here at all, since the concern about the moral status of actions seems to collapse into the issue of how luck can afflict the epistemic status of certain judgements that will be assessed retrospectively. That is, it collapses into the issue of how one can never, given the possibility of intrinsic luck, have the appropriate epistemic justification *in advance* to adequately rationally justify setting out on a project of this sort. Indeed, Williams is aware of this, noting that even if Gauguin's project were to be rationally justified in the way that he imagines, this does not mean that Gauguin will thereby have "any way of justifying himself to others, or at least to all others" (Williams 1981*b*, 23) and, plausibly, an ability to offer such a justification is constitutive of one's action enjoying a positive moral status.¹¹ So even if one could evade the problem of intrinsic luck at issue here and recover one's rational justification, this need not have any affect on the moral status of one's actions (i.e., his actions could be regarded as immoral regardless of whether he is successful in his project). Accordingly, unlike the examples that Nagel offers, there is no clear reason for thinking that the presence of luck can affect the moral status of one's actions.

Furthermore, even the underlying epistemological problem that Williams alludes to is unclear. For let us suppose, for the sake of argument, that the beliefs upon which Gauguin's decision is made (his beliefs about what his artistic abilities are, for example), are all true and enjoy an adequate positive epistemic status which suffices to ensure that his decision is based on knowledge. The problem is that insofar as Gauguin really does know that he has the great talent that he thinks he has, and insofar as we exclude in advance the possibility of extrinsic luck so that his project is genuinely tested, then it is hard to see how he could possibly fail in this enterprise. Wouldn't failure in this regard simply indicate that he lacked the relevant knowledge after all? Put another way, to have one's projects frustrated by intrinsic luck alone seems to imply that the beliefs upon which that project was based were epistemically faulty in some way, and an ascription of knowledge to the propositions in question rules out this possibility.

So what then is the epistemic problem that Williams is focussing upon here? I think that we get an idea of what he has in mind in the following passage:

[T]here might be grounds for saying that the person who was prepared to take the decision, and was in fact right, actually knew that he would succeed [...] But even if this is right for some cases, it does not help with the problems of retrospective justification. For the concept of knowledge here is itself applied retrospectively, and while there is nothing wrong with that, it does not enable the agent at the time of his decision to make any distinctions that he could not already make. As one might say, even if it did turn out in such a case that the agent did know, it was still luck, relative to the considerations available to him at the time [...] that he should turn out to have known. (Williams 1981*b*, 25-6).¹²

Reading between the lines here, one can take it that the "considerations [*that were*] available to

[*Gauguin*] at the time” he made his decision were those considerations that were reflectively accessible to him. Williams’ point therefore comes down to the claim that in terms of what the agent is able to know by reflection alone, it is matter of luck that the agent knows that the decision in question is the right decision. That is, that no matter how good the reflectively accessible grounds are upon which an agent bases his decision, it will still be possible for a substantial degree of intrinsic luck to intervene such that his beliefs, though well founded, were not knowledge after all (and thus that his decision was not rationally justified).

Notice, however, that the type of epistemic luck that is at issue here is not merely the luck that one’s beliefs are true, which is the sort of *veritic* luck that is at issue in, for example, epistemological debates about the style of counterexamples introduced by Edmund Gettier (1963). In these cases the worry is simply that one might have a true belief in such a way that one meets a certain kind of epistemic rubric (such as the classical tripartite demand for justification) and yet fail to have knowledge because one’s belief is only luckily true. For example, one might be looking at a stopped clock which one has every reason for thinking is working, and yet nevertheless form a true belief because one happens to look at the clock at one of the two times in the day when it is ‘telling’ the right time. Here the truth of one’s belief is lucky in that although one has a true belief in the actual world, in most near-by possible worlds where the relevant initial conditions are the same—in particular, where one forms one’s beliefs in the same manner as one does in the actual world—one will form a false belief. We thus get the following characterisation of veritic luck:

Veritic Luck

For all agents, φ , the truth of an agent’s belief in a contingent proposition, φ , is veritically lucky if, and only if, the agent’s belief that φ is true in the actual world, but false in most near-by possible worlds in which the belief is formed in the same manner as in the actual world.

Famously, the moral of Gettier is that no purely internalist theory of knowledge—and in particular the classical tripartite account that viewed knowledge as (internalistically) justified true belief—could be an acceptable theory. The reason for this is that purely internal epistemic conditions such as the justification condition, classically conceived, are insufficient to rule out the possibility of veritic luck. By an internal epistemic condition here we mean, of course, a condition which incorporates an ‘access’ requirement such that the agent has special access to the fact that this condition has obtained. The standard way of defining an internal condition is by saying that the facts which determine that the condition has obtained are facts that the agent can know by reflection alone,¹³ as we would expect to be the case when it comes to the classical justification condition which demands that the agent has (and is in a position to adduce) sufficient reasons in favour of his beliefs. Critically, however, no matter how good the supporting reasons one has for one’s beliefs are, this will be

consistent with one forming a true belief in a way that is veritically lucky. For example, no matter what grounds one might have for thinking that a clock is working, these grounds are consistent with it in fact having stopped, and thus consistent with one's belief about what the time is being 'Gettiered' by being, as it happens, true even though one gained that true belief luckily via a stopped clock.

The only way to deal with Gettier-style cases is thus to adduce an external epistemic condition—one that does not incorporate an internalist access requirement—which ensures that one's belief is not veritically lucky. One could, for example, demand that one's beliefs must be *safe*, where this means that in most near-by possible worlds in which one forms one's belief in the same way as in the actual world, one continues to form a true belief. Clearly, such a condition is going to be an external epistemic condition because the facts that determine whether or not one's beliefs are safe are not going to be facts that the agent can know by reflection alone. Moreover, the adoption of such a condition will obviously deal with the problem of veritic luck since it *defines* such luck out of existence.¹⁴

The type of luck that Williams has in mind is not of this standard veritic type, however, since if Gauguin's beliefs are in the market for being counted as knowledge at all, then they had better be safe. Williams' worry is not whether, as a matter of fact, Gauguin has formed his beliefs in such a way that they will track the truth across the relevant near-by possible worlds, but rather the more specific concern that Gauguin lacks reflectively accessible grounds which are sufficient to exclude the possibility that he is substantively wrong in his judgements about his ability. Gauguin's beliefs could be, as it happens, safe, without his beliefs enjoying the subjective epistemic assurance that his life-changing decision appears to demand.

Indeed, one's beliefs could be safe, and thus not veritically lucky, without one being able to offer *any* supporting reasons or considerations in favour of one's beliefs at all. Accordingly, the problem that Williams raises regarding the inadequacy of our reflectively accessible grounds need not gain a purchase on the issue of whether or not we have knowledge in such cases at all. This is the issue on which, famously, internalists and externalists in the theory of knowledge diverge. On the one hand, there are those who think that meeting an internal epistemic condition is essential to knowledge, while, on the other, there are those who think that meeting such a condition, though epistemically desirable, is inessential to knowledge. The debate thus tends to cluster around a set of cases where the two views come apart.

One such example is the case of the 'chicken sexer'. Here we asked to imagine an agent who, in virtue of being raised around chickens, has the highly reliable ability to tell male and female chicks apart. Crucially, however, our protagonist has false beliefs about how he is doing what he is doing (he thinks he is touching or seeing something distinctive, even though there is nothing distinctive for him to see or touch), and is not yet availed of his success ratio in this regard. Accordingly, he doesn't seem

to meet any internal epistemic condition at all since he has hardly any reflectively accessible supporting grounds that he can offer in favour of his beliefs. Nevertheless, since he really does have the ability in question those of an epistemically externalist persuasion will be inclined to maintain that he has done enough to be credited with knowledge of the sex of the chicks. His beliefs are, for example, safe (and thus ‘Gettier-proof’), since insofar as he is forming a true belief in this manner in the actual world, he will be likewise forming true beliefs in this manner in most near-by possible worlds as well (and thus his belief will not be susceptible to Gettier-style counterexamples). Those impressed by epistemically internalist intuitions will balk at this suggestion, however, arguing that agents need to do more than merely meet external epistemic conditions if they are to know. In particular, they will claim that the lack of reflectively accessible grounds available to the subject in this case precludes him from having knowledge.¹⁵

Indeed, one could put the point that the epistemic internalist is making here in terms of luck. Linda Zagzebski writes, for example:

The dispute between externalists and internalists looms large mostly because of ambivalence over the place of luck in normative theory. Theorists who resist the idea that knowledge [...] is vulnerable to luck are pulled in the direction of internalism [...] Externalists are more sanguine about luck. [...] There is lots of room for luck in externalist theories since the conditions that make it the case that the knower is in a state of knowledge are independent of her conscious access. (Zagzebski 1996, 39)

This isn’t quite right, of course, since even epistemic externalists are concerned to eliminate veritic luck from their theories, but it does seem that Zagzebski is on to something here, in that epistemic externalists are sanguine about a certain kind of luck in a way that epistemic internalists are not. I think the kind of luck that is at issue here is what we might term a *reflective* form of veritic luck. The issue is still the luck that infects the truth of one’s beliefs, but whereas when it comes to veritic luck the possible worlds are ordered in the usual way in terms of the facts of the situation, when it comes to reflective luck the ordering is determined by what the agent is able to know by reflection alone.

In order to see this distinction, consider two chicken sexers, one understood as we described him above as being ‘naïve’ in the relevant sense (since he knows next to nothing by reflection alone that is relevant to the subject matter in hand), and a counterpart ‘enlightened’ chicken sexer who is exactly alike the naïve chicken sexer except that he knows a great deal by reflection alone regarding the subject matter in question (in particular, he has good reflectively accessible grounds which indicate not just how he is able to tell the chicks apart, but also that his ability is reliable). Neither of these chicken sexers are forming beliefs that are veritically lucky in the sense outlined above, and thus neither of them are forming true beliefs that could be ‘Gettiered’. Nevertheless, there does seem to be something lucky about the naïve chicken sexer’s true beliefs that is not lucky about his enlightened counterpart’s beliefs, and this relates to the fact that, in terms of what the naïve chicken sexer knows by reflection

alone, there is no reason for thinking that the near-by possible worlds will be worlds where he continues to form a true belief. Indeed, if we only take what he can know by reflection alone into account when ordering the possible worlds, then just about *any* possible world can count as near-by in the relevant respect, including a wide class of possible worlds where he no longer has the ability in question and so forms false beliefs about the sex of the chicks as a result. This won't be the case for his enlightened counterpart since his excellent reflective knowledge will ensure that the ordering of the possible worlds will be roughly the same as the 'objective' ordering in terms of the facts of the situation. We thus have the following understanding of reflective luck:

Reflective Luck

For all agents, φ , the truth of an agent's belief in a contingent proposition, φ , is reflectively lucky if, and only if, the agent's belief that φ is true in the actual world, but, *given only what the agent is able to know by reflection alone*, false in most near-by possible worlds in which the belief is formed in the same manner as in the actual world.

Whereas both the naïve and the enlightened chicken sexers will be forming beliefs in ways that are not veritically lucky, only the enlightened chicken sexer will be forming beliefs in ways that are not reflectively lucky either.¹⁶

In focussing upon a specifically reflective form of epistemic luck, it seems to be this sort of luck that Williams has in mind. The epistemic worry that Williams is giving expression appears to concern how no matter what the pedigree is of the reflectively accessible grounds that an agent like Gauguin adduces when making his life-changing decision to abandon his family, it will remain that his belief that he will succeed is reflectively lucky. Interestingly, on the issue of whether or not this reflective luck should undermine the possibility of his knowing that he will succeed, Williams is ambivalent. He notes that there "might be grounds for saying that person who was prepared to take the decision, and was in fact right, actually knew that he would succeed", and then further remarks that "even if it did turn out in such a case that the agent did know, it was still luck, relative to the considerations available to him at the time [...] that he should turn out to have known" (Williams 1981*b*, 25-6). Williams thus does not definitively take sides on the epistemic externalism/internalism distinction. Nevertheless, it should be clear from how Williams puts matters in this way that it must be reflective luck that is at issue, since the point seems to be that from the agent's point of view at the time that he made the decision he has insufficient reflectively accessible grounds for thinking that what he believes (i.e., that he will succeed) is true even though, provided he does indeed meet the relevant external conditions for knowledge, then he may well have knowledge of what he believes (at least by the lights of an externalist account of knowledge). The ambivalence thus relates to the possibility of taking two different viewpoints as regards the knowledge in question—one which is 'internal' and only takes into account what the agent has reflective access to; and one which is 'external' and takes other relevant

facts (such as whether or not his belief is safe) into account. Clearly these two viewpoints will roughly map onto our distinction between reflective and veritic luck, in that in the former case the possible worlds are ordered ‘internally’ in terms of only what the agent can know by reflection alone rather than in terms of the actual facts of the situation. Williams’ point is thus that whether or not the agent in the case that he imagines knows that he will succeed is subject to reflective epistemic luck.

There is clearly something right about Williams’ claim here, but it will not do as it currently stands. The reason for this is that the kind of reflectively accessible grounds that the agent has *will* suffice to eliminate reflective luck in any substantive degree from his knowledge, at least insofar as reflective luck is ever eliminable. Compare, for example, Gauguin’s belief that he will succeed with that of the enlightened chicken sexer. Just as the enlightened chicken sexer, unlike his naïve counterpart, evades the problem of reflective luck by having excellent reflectively accessible grounds in favour of his belief, so the same should be true of Gauguin in the example that Williams offers. If he does really have the supporting grounds in question, then the near-by possible worlds on the ‘reflective’ ordering should be such that in most of them he continues to form a true belief (in the same way as in the actual world) about whether or not he succeeds. Of course, his beliefs might not be knowledge because they are not safe, but that is to take them out of the market for knowledge altogether, by the lights of an externalist *or* an internalist theory of knowledge. So what then is the problem that Williams is giving expression to here?

I think that the answer lies in the fact that we only count the enlightened chicken sexer as eliminating reflective luck from his beliefs because we are implicitly bracketing the sceptical possibility that his reflectively accessible grounds might bear no relation to the world whatsoever. Imagine, for example, that the enlightened agent in the chicken sexer example was, unbeknownst to him, a brain-in-a-vat being ‘fed’ his experiences by neuroscientists. Clearly in this case the reflectively accessible grounds that he has in favour of his beliefs are wholly unreliable since they are in no way indicative of the truth. Provided that we do not ‘factor-out’ such sceptical scenarios, however, then the ordering of the possible worlds that only takes into account the agent’s reflective knowledge will include a wide class of near-by possible worlds in which the agent is the victim of a sceptical scenario since, as the sceptic famously points out, we do not have any adequate reflectively accessible grounds for thinking that we are not the victim of a sceptical hypothesis. Accordingly, his beliefs cannot help but be reflectively lucky to a substantive degree. The same goes for Gauguin. Unless we exclude the sceptical scenarios from our calculations of how the possible worlds should be ordered given only what he knows by reflection alone, then there will be a wide class of near-by possible worlds in which he forms his belief in the same way as in the actual world and yet forms a false belief as a result because he is the victim of a sceptical error-possibility (e.g., he is a brain-in-a-vat who has been tricked into

thinking that he has the abilities that he took himself to have).

It is not my intention to try to resolve this sceptical problem here. The crucial point for our purposes is that the problem that Williams raises as regards Gauguin, like the original problem that we saw the epistemological internalist about knowledge raising about the naïve chicken sexer above, makes no mention of the problem of scepticism at all. Accordingly, it is odd to discover that one can only make sense of that problem by understanding it in the light of sceptical concerns. It seems, then, that Williams not only (perhaps intentionally) fails to distinguish between the problem of moral luck and the analogue problem of epistemic luck, but that he also (seemingly unintentionally) fails to recognise that the problem of epistemic luck that he raises (insofar as we can make sense of it at all) is really the very specific issue about epistemic luck highlighted by the sceptic.¹⁷

4. NAGEL ON EPISTEMIC LUCK AND SCEPTICISM

A closer inspection of Nagel's paper on moral luck further reveals how his argument for the existence of moral luck, like Williams', in fact implicitly draws on the problem of epistemic luck and, in particular, the specific problem of reflective epistemic luck identified by the sceptic. This is significant, because if the motivation that Nagel presents us for taking the problem of moral luck seriously—a problem which, I have argued, Nagel offers us no good grounds for thinking is genuine—makes appeal to the putatively analogous sceptical problem of reflective epistemic luck—a problem which I claim *is* genuine—then it is hardly surprising that we find ourselves attracted to the intuitions that Nagel offers.

For example, at one point in the paper Nagel argues that the problem of moral luck

[...] resembles the situation in another area of philosophy, the theory of knowledge. There too conditions which seem perfectly natural, and which grow out of the ordinary procedures for challenging and defending claims to knowledge, threaten to undermine all such claims if consistently applied. Most sceptical arguments have this quality: they do not depend on the imposition of arbitrarily stringent standards of knowledge, arrived at by misunderstanding, but appear to grow inevitably from the consistent application of ordinary standards. There is a substantive parallel as well, for epistemological scepticism arises from consideration of the respects in which our beliefs and their relation to reality depends on factors beyond our control. External and internal causes produce our beliefs. We may subject these processes to scrutiny in an effort to avoid error, but our conclusions at the next level also result, in part, from influences which we do not control directly. The same will be true no matter how far we carry the investigation. Our beliefs are, ultimately, due to factors outside of our control, and the impossibility of encompassing those factors without being at the mercy of others leads us to doubt whether we know anything. It looks as though if any of our beliefs are true, it is pure biological luck rather than knowledge. (Nagel 1979, 26-7)

Two points are primarily significant here. The first is the dialectical observation that the remarks in this passage are being used to *motivate* the examples of moral luck that are subsequently offered, which implies that it is the phenomenon of epistemic luck that possesses the greater intuitive force. Moreover,

unlike Williams, Nagel is in this quotation explicitly identifying the problem of reflective epistemic luck with the specific version of that problem raised by the sceptic. Thus, it is not the problem of reflective epistemic luck *simpliciter* that is meant to be motivating his further (and as we saw above, erroneous) contentions regarding moral luck, but rather the particular sceptical use of this problem.

The second point to note is that the initial focus in this passage is on knowledge *claims*, as if what one knows one can, at least typically, properly claim to know. To put the emphasis on knowledge claims in this way is to thereby focus on knowledge of a very specific sort—knowledge where the agent has met an internal epistemic condition. The reason for this is that the kind of ‘brute’ knowledge allowed by the externalist where the agent has no (or hardly any) reflectively accessible grounds in support of his beliefs will not be of a sort to allow the agent to properly claim that knowledge (think, for example, of the naïve chicken sexer in this respect). This is because a claim to know is only conversationally appropriate if (*inter alia*) the agent concerned can offer adequate supporting grounds and brute knowers precisely lack such supporting grounds for their beliefs. All parties to the epistemological externalism/internalism dispute should thus be willing to grant that, *ceteris paribus*, meeting an internal epistemic condition is essential if one is to be in a position to properly claim knowledge. That Nagel views the problematic form of epistemic luck as undermining the propriety of knowledge claims thus indicates that it is a specifically reflective variety of epistemic luck that he has in mind.

Indeed, Nagel’s remarks about “lack of control” and the “internal” and “external” determinants of belief in this passage would seem to suggest that there are factors that are relevant to the epistemic status of an agent’s beliefs but which are not reflectively knowable by the agent. That is, that the factors relevant to knowledge that we are able to control—the ‘internal’ and thus, we might suppose, reflectively knowable factors—will not suffice to determine whether or not we do in fact have knowledge. Instead, external factors (and thus non-reflectively knowable) factors will also be relevant. Nagel’s point thus seems to be on a par with Williams’, in that he is contending that, in terms of the ‘internal’ reflectively accessible grounds possessed by the agent alone, it is a matter of reflective luck that his belief is true.

This interpretation of Nagel is further confirmed once one considers the famous example that Nagel mentions in a footnote to back-up his points about epistemic luck where he notes that “the Nobel Prize is not awarded to people who turn out to be wrong, no matter how brilliant their reasoning” (Nagel 1979, 183-4). The implication of this remark is that what the agent is in control of extends only so far as meeting the relevant internal conditions—such as ensuring that his reasoning is as impeccable as it can be—but that what ultimately determines knowledge (and thus enables one to be in the market for a Nobel Prize) goes beyond this to implicate external factors such as (primarily) whether or not the

belief in question is true.

In later works, where Nagel is explicitly dealing with the problem of scepticism rather than the problem of moral luck, he is much clearer about what this issue regarding reflective epistemic luck amounts to. He argues that objectivity involves attaining a completely impartial view of reality, one that is not tainted by any particular perspective. We aspire, he contends, to “get outside of ourselves”, and thereby achieve the impossible task of being able to “view the world from nowhere from within it” (Nagel 1986, 76). We realise that the initial appearances present to a viewpoint can be unreliable guides to reality and therefore seek to modify our ‘subjective’ view with a more ‘objective’ perspective that is tempered by reason and reflection. As Nagel points out, however, the trouble with this approach is that

[...] if initial appearances are not in themselves reliable guides to reality, [*then*] why should the products of detached reflection be any different? Why aren't they [...] equally doubtful [...]? [...] The same ideas that make the pursuit of objectivity seem necessary for knowledge make both objectivity and knowledge seem, on reflection, unattainable. (Nagel 1986, 76)

Of course, what Nagel means by “perspective” here essentially involves what the agent is able to know by reflection alone—i.e., our reflectively accessible grounds for believing what we do—which means that the problem he poses is that our demand for objectivity imposes a requirement that, in the best case at least, what we are able to know by reflection alone will entail that the world is the way we take it to be. Such a requirement cannot be met, however, since our reflectively accessible grounds are always going to be consistent with the truth of a sceptical scenario. Nagel thus draws the pessimistic conclusion that the problem of scepticism

[...] has no solution, but to recognise that is to come as near as we can to living in the light of truth. (Nagel 1986, 231)

We are therefore back to the problem of reflective epistemic luck posed by the sceptic.

It is important to note that if Nagel understands the problem of scepticism as being the problem of the possibility of knowledge *simpliciter*, then this conclusion goes much farther than that advanced by Williams. As we saw above, Williams merely presents a problem for our ordinary understanding of knowledge, in that we seem to have it in such a way that is consistent with a degree of reflective epistemic luck. Nagel, however, is here drawing the stronger sceptical conclusion from this observation—that since knowledge is incompatible with luck, so we must lack the knowledge we take ourselves to have after all. In effect, Nagel is endorsing a form of epistemological internalism which demands the complete elimination of reflective luck and concluding, since such a complete elimination is impossible, that the genuine possession of knowledge is impossible also.¹⁸

In any case, what is important to the present discussion is that it is a certain conception of the sceptical problem that is motivating Nagel's general concerns about epistemic luck which, in turn, are

being used to motivate the putative problem of moral luck. Provided we fix a suitable conception of luck throughout our consideration of these problems, however, then it becomes apparent that the genuine sceptical problem of reflective epistemic luck does not in itself give us any cause to treat the problem of moral luck that Nagel identifies as being similarly genuine. We can thus diagnose the superficial attraction of Nagel's position on moral luck as being in part due to the equivocation identified earlier in section two, and also in part due to the fact that he appeals to the genuine and apparently analogous sceptical problem of reflective luck in order to get us to take the illusory problem of moral luck seriously. Being clear about what luck is, and the role that it plays in discussions of epistemic and moral luck, thus enables us to both undermine Nagel's account of moral luck and also tell a powerful diagnostic story about why we were ever taken in by this account in the first place.

5. CONCLUDING REMARKS

Ultimately, both Williams and Nagel implicitly motivate their examples regarding moral luck—examples which we have found to be on reflection inconclusive—via an appeal to the sceptical problem of reflective epistemic luck. The immediate import of this observation is that the problem of moral luck identified by Williams and Nagel, insofar as they identify a problem at all, is not what they take the problem to be. The secondary import of this observation is that the epistemological problem of scepticism poses a general existential difficulty that has not only an abstract impact on our doxastic practices but also impacts on us in a more concrete fashion by undermining our ability to legitimate the life-changing courses of actions that we opt for. This highlights the sense in which the problem of scepticism is an ethical problem, in the broad sense of that term, a conclusion which, while familiar to the classical sceptics (such as the Pyrrhonians), is not so familiar to the contemporary mind-set.

While making this connection between the broadly ethical problem raised by Williams and Nagel and the problem of scepticism elevates the philosophical importance of the latter somewhat, it also by that very fact undermines the philosophical importance of the former. For while we might wish a resolution to the sceptical problem all the more now that its ethical ramifications are made clear, we also recognise that a great number of debates have to take place in the absence of a solution to the problem of epistemological scepticism. In this sense, then, the sceptical problem posed here is on a par with that posed in many other areas of philosophy.¹⁹

REFERENCES

- Bennett, J. (1995). *The Act Itself*, Oxford University Press, Oxford.
- Brandom, R. (1998). 'Insights and Blindspots of Reliabilism', *The Monist* 81, 371-92.
- Browne, B. (1992). 'A Solution to the Problem of Moral Luck', *The Philosophical Quarterly* 42, 345-56.
- Card, C. (1990). 'Gender and Moral Luck', *Identity, Character, and Morality*, (eds.) O. Flanagan & A. Rorty, 199-218, MIT Press, Cambridge, Massachusetts.
- Dancy, J. (1985). *Introduction to Contemporary Epistemology*, Basil Blackwell, Oxford.
- Foley, R. (1987). *A Theory of Epistemic Rationality*, Harvard University Press, Cambridge, Massachusetts.
- Gettier, E. (1963). 'Is Justified True Belief Knowledge?', *Analysis* 23, 121-3.
- Greco, J. (1995). 'A Second Paradox Concerning Responsibility and Luck', *Metaphilosophy* 26, 81-96.
- Greco, J. (2004). 'Truth in Light of Virtue', *manuscript*.
- Jensen, H. (1984). 'Morality and Luck', *Philosophy* 59, 323-30.
- Kornblith, H. (2001). *Epistemology: Internalism and Externalism* (ed.), Basil Blackwell, Oxford.
- Latus, A. (2000). 'Moral and Epistemic Luck', *Journal of Philosophical Research* 25, 149-72.
- Latus, A. (2003). 'Constitutive Luck', *Metaphilosophy* 34, 460-75.
- Lewis, D. (1996). 'Elusive Knowledge', *Australasian Journal of Philosophy* 74, 549-67.
- Moore, A. W. (1990). 'A Kantian View of Moral Luck', *Philosophy* 65, 297-321.
- Nagel, T. (1976). 'Moral Luck', *Proceedings of the Aristotelian Society* (supplementary volume) 76, 136-50.
- Nagel, T. (1979). 'Moral Luck', *Mortal Questions*, 24-38, Cambridge University Press, Cambridge, UK.
- Nagel, T. (1986). *The View from Nowhere*, Oxford University Press, Oxford.
- Nozick, R. (1981). *Philosophical Explanations*, Oxford University Press, Oxford.
- Pritchard, D. H. (2002a). 'Recent Work on Radical Skepticism', *American Philosophical Quarterly* 39, 215-57.
- Pritchard, D. H. (2002b). 'Resurrecting the Moorean Response to the Sceptic', *International Journal of Philosophical Studies* 10, 283-307.
- Pritchard, D. H. (2003). 'Virtue Epistemology and Epistemic Luck', *Metaphilosophy* 34, 106-30.
- Pritchard, D. H. (2004). 'Epistemic Luck', *Journal of Philosophical Research* 29, 193-222.
- Pritchard, D. H. (2005a). *Epistemic Luck*, Oxford University Press, Oxford.
- Pritchard, D. H. (2005b). 'Scepticism, Epistemic Luck and Epistemic Angst', *Australasian Journal of Philosophy* 83.
- Pritchard, D. H. (forthcominga). 'Scepticism as a Problem about Luck', *The Oxford Handbook of Scepticism*, (ed.) J. Greco, Oxford University Press, Oxford.
- Pritchard, D. H. (forthcomingb). 'Externalism, Skepticism and Luck', *Internalism and Externalism in Semantics and Epistemology*, (ed.) S. Goldberg, Oxford University Press, Oxford.
- Pritchard, D. H., & Smith, M. (2004). 'The Psychology and Philosophy of Luck', *New Ideas in Psychology* 22, 1-28.
- Pryor, J. (2001). 'Highlights of Recent Epistemology', *The British Journal for the Philosophy of Science* 52, 95-124.
- Rescher, N. (1993). 'Moral Luck', *Moral Luck*, (ed.) D. Statman, 141-66, State University of New York Press, Albany, New York.
- Richards, N. (1993). 'Luck and Desert', reprinted in *Moral Luck*, (ed.) D. Statman, 167-80, State University of New York Press, Albany, New York.
- Sosa, E. (1999). 'How to Defeat Opposition to Moore', *Philosophical Perspectives* 13, 141-54.
- Statman, D. (1991). 'Moral and Epistemic Luck', *Ratio* 4 (new series), 146-56.
- Statman, D. (1993a). 'Introduction', *Moral Luck*, (ed.) D. Statman, 1-34, State University of New York

- Press, Albany, New York.
- Statman, D. (1993*b*). *Moral Luck* (ed.), State University of New York Press, Albany, New York.
- Thomson, J. J. (1993). ‘Morality and Bad Luck’, reprinted in *Moral Luck*, (ed.) D. Statman, 195-215, State University of New York Press, Albany, New York.
- Williams, B. (1976). ‘Moral Luck’, *Proceedings of the Aristotelian Society* 50 (supplementary volume), 115-35.
- Williams, B. (1981*a*). *Moral Luck*, Cambridge University Press, Cambridge, UK.
- Williams, B. (1981*b*). ‘Moral Luck’, *Moral Luck*, 20-39, Cambridge University Press, Cambridge, UK.
- Williams, B. (1993). ‘Postscript’, *Moral Luck*, (ed.) D. Statman, 251-8, State University of New York Press, Albany, New York.
- Zagzebski, L. (1996). *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*, Cambridge University Press, Cambridge, UK.
- Zimmerman, M. J. (1987). ‘Luck and Moral Responsibility’, *Ethics* 97, 374-86.

NOTES

- ¹ These papers were subsequently reprinted (in a slightly altered form) as Nagel (1979) and Williams (1981*b*), and it is these versions of the papers—which have each been reprinted since in a number of different anthologies—that I shall focus upon here.
- ² For some of the key discussions of the exchange between Williams and Nagel, see Statman (1993*b*).
- ³ Following Nagel and Williams, a number of other authors have made this connection between luck and lack of control, making it what is, perhaps, the most common philosophical understanding of the notion. Consider these passages from Statman (1991), Zimmerman (1993) and Greco (1995):
- “Let us start by explaining what we usually mean by the term ‘luck’. Good luck occurs when something good happens to an agent P, its occurrence being beyond P’s control. Similarly, bad luck occurs when something bad happens to an agent P, its occurrence being beyond his control.” (Statman 1991, 146)
- “[S]omething which occurs as a matter of luck with respect to someone P is something which occurs beyond P’s control.” (Zimmerman 1993, 231)
- “[T]o say that something occurs as a matter of luck is just to say that it is not under my control.” (Greco 1995, 83)
- See also Latus (2000). Latus (2003, 446) also offers the following examples of this type of view about luck:
- “‘As a matter of luck’ here means: in a way that is beyond our control.” (Moore 1990, 301)
- “By ‘luck’ I mean factors, good or bad, beyond the control of the affected agent.” (Card 1990, 199)
- ⁴ A critical survey of the philosophical literature on luck along with a detailed defence of this characterisation of luck in the light of that survey is offered in Pritchard (2004; 2005*a*) and Pritchard & Smith (2003).
- ⁵ Nagel actually begins the article by simply talking about a morally lucky reckless driver, without specifying that the recklessness in question has anything to do with drink. Later on, however, he specifically mentions an example of a reckless driver who is drunk. For convenience, rather than take these as two different examples, I will simply regard them as one example where the recklessness is in both cases due to drink.
- ⁶ For examples of responses along these general lines, see Jensen (1984), Rescher (1993), Richards (1993), Thomson (1993) and Bennett (1995). For a discussion specifically on the relationship between moral luck and punishment, see Browne (1992).
- ⁷ Nagel also distinguishes these two types of luck in turn from what he terms “causal” and “constitutive” luck. These are, respectively, “luck in how one is determined by antecedent circumstances” and the luck involved in a person’s having the “inclinations, capacities and temperament” that he does (Nagel 1979, 27). For the sake of brevity I will not be extending this discussion to cover examples of these kinds of luck here, although I think that it can be applied. In any case, much of the focus of Nagel’s article is on resultant and circumstantial luck. For an interesting recent discussion that is explicitly focussed on constitutive luck, see Latus (2003).
- ⁸ Both Rescher (1993) and Richards (1993) argue that all that cases of moral luck highlight is merely an epistemic lack on our part to what the appropriate moral judgement of someone’s actions should be. That is, as Statman (1993*a*, 17) puts it, “luck does not affect one’s *deserts* but only our *knowledge* of them.” See also Zimmerman (1987).
- ⁹ Nagel begins with cases like the drunk driver example, and then moves onto examples of circumstantial luck like the ‘German expatriate’ example. Greco (2004) argues for a similar conclusion.
- ¹⁰ Williams wouldn’t put the point in these terms, of course, since his ultimate goal is to deflate the three-way distinction between the moral, the ethical, and the practically rational. For more on this point, see Williams (1981*a*; 1993).

¹¹ Elsewhere in the paper Williams describes the Kantian view that he is opposed to as demanding that in terms of “the agent’s reflective assessment of his own actions [...] it cannot be a matter of luck whether he was justified in doing what he did.” (Williams 1981*b*, 23).

¹² The reader should note that I’ve deleted a *caveat* from this quotation that might be thought to be significant. This is where Williams notes that the successful agent might have known that he would have succeeded “however subjectively uncertain he may have been” at the time. I’ve removed this phrase since it ought not to be relevant to the case in hand. If Williams is right, then it should be possible to construct a Gauguin-type example where it is explicitly stipulated that Gauguin is subjectively certain of the correctness of his decision. The *caveat* thus only serves to add a complication to the proceedings that is irrelevant to the main thrust of the argument.

¹³ This is, for example, roughly the way that Pryor (2001) defines the notion of an internal epistemic condition in his recent survey article in epistemology.

¹⁴ For more on the safety principle, see Sosa (1999) and Pritchard (2002*b*). This isn’t, of course, the only way to eliminate veritic luck from one’s beliefs since incorporating other principles—such as Nozick’s (1981) ‘sensitivity’ principle—into one’s theory of knowledge would also do the trick. For more on the point that safety-based views are, essentially, anti-veritic luck epistemologies, see Pritchard (2003; 2004; 2005*a*, chapter six).

¹⁵ For more discussion of the chicken sexer example, see Foley (1987, 168-9), Lewis (1996), Zagzebski (1996, §2.1 & §4.1), and Brandom (1998). For more on the epistemological externalist/internalist distinction in general, see the papers collected in the excellent anthology edited by Kornblith (2001).

¹⁶ For more on this distinction between veritic and reflective luck—including discussion of cases like the chicken sexer example and their relationship to the debate about epistemic luck—see Pritchard (2003; 2004; 2005*a*, chapter six; 2005*b*; *forthcominga*; *forthcomingb*).

¹⁷ For further discussion of the point that scepticism can be understood as a problem regarding the ineliminability of reflective epistemic luck, see Pritchard (2004; 2005*b*; *forthcominga*).

¹⁸ For more on the contemporary debate regarding radical scepticism, and the place of Nagel’s work within it, see Pritchard (2002*a*).

¹⁹ Thanks to Michael Brady, Antony Duff, and John Greco for discussion of material in this paper, and to The Leverhulme Trust for the award of a two-year Special Research Fellowship which has enabled me to conduct work in this area.