

---

## Ethics versus Luck?

### *The myriad forms of luck*

Is 'moral' incompatible with luck? The attempt to shore up the ethically right against the tidal waves of ill-chance is often thought central to ethical thought. If it fails, is ethics as a whole somehow at risk? Ethics is commonly assumed to be the one realm in which luck does not intrude. 'While one can be lucky in one's business, in one's married life, and in one's health, one cannot, so it is commonly assumed, be subject to luck as far as one's moral worth is concerned.'<sup>1</sup>

The imperviousness of ethics to luck is central to the 'appeal' of ethics. If ethics is not impervious to luck, the notion of the moral enterprise as a universal good suffers doubly. First, justice as equality is threatened: it is not open to some – those tainted by the particular kinds of ill-luck that I shall come to – to be as moral as those who are not so affected. Secondly, justice as desert is infringed: 'if morality depends on luck, then at least sometimes people are judged morally for things that are beyond their control.'<sup>2</sup>

Typically we maintain two incompatible standards towards right actions and good character, and the tension between these polarities creates the paradox of moral luck. In practice we regard actions as right or wrong, and moral character as good or bad, partly according to what happens as a result of the agent's decision. That is, we make responsibility hinge to some extent on things outside the agent's control. Yet at the same time we think that people should not be held responsible for matters beyond their control.

This tension underpins Kant's famous assertion that only the good will is securely good, and that its goodness is impervious to ill-luck in how things actually turn out.

## 2 *Ethics versus Luck?*

---

Even if it should happen that by a particularly unfortunate fate or by the niggardly provision of a stepmotherly nature, this will should be wholly lacking in power to accomplish its purpose, if even the greatest effort should not avail it to achieve anything of its end, and if there remained only the good will (not as a mere wish, but as the summoning of all the means in our power), it would sparkle like a jewel in its own right, as something that had full worth in itself.<sup>3</sup>

But is moral luck only a problem for a Kantian? Perhaps only one particular school of ethics is vulnerable to moral luck arguments, but ethics as a whole is not. That would be a reassuring conclusion for clinicians and philosophers alike. However, there is an increasing consensus that the moral luck problem is larger than that. Martha Nussbaum, for example, has claimed that not only Kant but also Plato was motivated by a concern to minimize the effects of chance on moral character and the rightness of ethical choices.<sup>4</sup> Although ethical consequentialism, including utilitarianism, appears not to have a problem with moral luck, it actually does, as I shall argue in chapter 3. Indeed, Bernard Williams, who originated the term ‘moral luck’, has always maintained that moral luck is a problem for ethics in general, not merely for Kantian ethics in particular:

[Non-Kantians] may be disposed to think, so far as morality is concerned, that all that is in question is the pure Kantian conception, and that conception merely represents an obsessional exaggeration. But it is not merely that, nor is the Kantian attempt to escape luck an arbitrary enterprise. The attempt is so intimate to our notion of morality, in fact, that its failure may rather make us consider whether we should not give up that notion altogether.<sup>5</sup>

If the paradox of moral luck is a genuine paradox, and if its applicability is really so vast, then the scope of our examination needs to be narrowed. In this book I want to examine whether the concepts of luck and risk can cast any light on problems in medical ethics, and, conversely, whether grounding the concepts of luck and risk in practical ethics might help us to better understand and perhaps resolve the paradox. The moral luck debate, begun by Bernard Williams and Thomas Nagel in 1976,<sup>6</sup> has become reasonably familiar to many philosophers, but is still comparatively little known in medical ethics proper.<sup>7</sup> This is unfortunate, since the concept of moral luck can offer a good deal of help to clinical ethicists, health care professionals and students, managers and others involved in making decisions in med-

ical ethics. The basic question to ask is simple: what happens if things go wrong?

In the first three chapters, I erect the theoretical framework for what I hope will actually be a more sophisticated discussion than might be suggested by the basic question of what happens when things beyond an agent's control do go wrong. I begin in this chapter with a systematic exposition of the moral luck debate between Williams and Nagel, considering also the ramifications of the debate in further articles on the same subject by other writers.<sup>8</sup> I explore Williams's own further reflections on the significant debate which he set in train.<sup>9</sup> Whether the paradox is genuine, and, if it is a true paradox, the range of its effects on ethics, are also considered.

The second chapter begins by examining moral luck and virtue ethics. The moral luck debate was extended by Martha Nussbaum in *The Fragility of Goodness*, which concluded that only an Aristotelian, virtue-centred approach can overcome the paradox. The current popularity of virtue ethics in medicine<sup>10</sup> perhaps rests on its concern with what sort of doctor it is good to be, rather than what action is prescribed in disconnected individual decisions. This chapter will counter the accepted wisdom by arguing that a classical, virtue-centred approach does not in fact get round the problem of risk and luck in medical decision-making. Using Nussbaum's critique of Kant as a bridge, I also analyse the deontological, Kantian sources of the claim that ethics cannot allow itself to be undermined by chance. Nussbaum's rather conventional interpretation of Kant presents him as particularly prone to the moral luck paradox, in what she views as an attempt to create a risk-free ethics in which there remains only the good will. I shall argue instead that Kant wants conscience to have an impact on the world, and that a Kantian approach to dilemmas in the medical world, such as informed consent, actually offers a promising way forward, by limiting what it is that we are actually responsible for but interpreting that responsibility strictly.

In chapter 3 I turn from virtue ethics and deontology to ethical consequentialism, and particularly to utilitarianism. In contrast to deontology, utilitarianism makes the moral agent responsible for too much, and is particularly vulnerable to the paradox of moral luck. This is true even if the agent is responsible not for actual but for potential consequences. Three important concepts associated with moral luck are explored in this chapter: risk, probability and rationality. Chapters 2 and 3 are intended for the reader with philosophical interests, although they do not require any substantial philosophical background.

#### 4 *Ethics versus Luck?*

---

The reader whose interests are primarily clinical may prefer to press on to the applied chapters and to the conclusion in chapter 10. However, it might be worth a brief stopover at the section on remorse and regret at the end of chapter 3.

The applied, practical chapters of the book, chapters 4 through 9, concern, respectively:

- risk and luck in informed consent to treatment
- moral luck in decisions about withholding life-sustaining treatment
- moral luck in the allocation of scarce medical resources
- risk and luck in reproductive ethics, with particular focus on contract motherhood ('surrogacy'), and in cloning
- risk in psychiatry
- genetics and luck in character.

In several of these chapters, particularly those on allocation of resources, reproductive ethics and research ethics, I deliberately take a 'global' focus. I do not confine resource allocation to decisions within one jurisdiction; I consider the possible exploitation of third-world women as egg donors in stem cell technologies; and I ask questions about what risks it is permissible for first-world researchers to impose on third-world populations. The last chapter, chapter 10, considers the wider ramifications for global ethics of value pluralism, which Williams insists is the logical corollary of accepting that luck does fatally undermine ethics. In a final synthesis, I argue that radical value pluralism and its associated strands in global ethics can be avoided, and that we can look to similar debates in feminist theory about difference and identity to help us find a way out of the dilemma.

#### *A preliminary typology of luck*

In this first chapter I summarize the Williams–Nagel debate, and ask whether the moral luck paradox is genuine. Before I begin, however, I need to distinguish among various types of luck. The moral luck paradox has generally been construed as being primarily about *luck in outcomes or consequences*, with the question of what kinds of consequences we are responsible for. Within outcome luck, one might further distinguish (as does Ronald Dworkin) between *option luck* and '*brute*' *luck*. As Dworkin puts it:

Option luck is a matter of how deliberate and calculated gambles turn out – whether someone gains or loses through accepting an isolated risk he or she should have anticipated and might have declined. Brute luck is a matter of how risks fall out that are not in that sense deliberate gambles. If I buy a stock on the exchange that rises, then my option luck is good. If I am hit by a falling meteorite whose course could not have been predicted, then my bad luck is brute (even though I could have moved just before it struck if I had any reason to know where it would strike).<sup>11</sup>

The moral language of ‘should’ (‘should have anticipated and might have declined’) can apply only to option luck, which involves *choice* as opposed to mere *chance* (or, more properly, involves choice as well as chance). Put another way, I am responsible for the decisions which are involved in option luck, even though I am not solely responsible for the consequences. If my stock rises, I might congratulate myself on my financial acumen; if it falls, I might blame myself for choosing badly, for not getting out of a bear market sooner, or for indulging in the high-flown gambling which is the stock market in the first place. The extent to which I am responsible is precisely what we want to determine, and what gives rise to the paradox of moral luck. But although I might rue the moment I unwittingly stood in the path of the meteorite (assuming I somehow survive), it would be unreasonable of me to blame myself for my brute ill-luck. Later I will introduce a somewhat similar distinction made by Bernard Williams between regret and remorse, and a parallel distinction which he draws between luck intrinsic to the agent’s ‘project’ and luck extrinsic to it.

Ethical issues are not present in the second subcategory of outcome luck – brute luck – but neither are they confined to outcome luck alone. Outcome luck focuses our attention on results, but we must also look to luck in causal factors. If I give £10 to Oxfam as an expression of my belief that it is ethically desirable to help feed the hungry, I do not regard my decision as having been proved unethical by a flash flood which stops the supplies that I have donated from getting through.<sup>12</sup> However, I am of course lucky to have £10 to give, so that my generosity is somehow slightly less than mine alone: it is partly also a matter of my good fortune in living in a comparatively wealthy country and being in work. (Of course, some people live in a first-world country, hold paid employment, and do not contribute to charity, so my generosity is partly my own.) This can be seen as a form of *circumstantial luck*, or *luck in antecedent circumstances*, to use Thomas Nagel’s terminology.<sup>13</sup> The sixth chapter of this book, on fair

distribution of scarce medical resources, will concern luck in antecedent circumstances as well as outcome luck.

I am also lucky to live under a watered-down version of Christianity which does not demand that I sell all I have and give my goods to the poor. I can be accounted generous only because the moral standards set by my society are actually rather lax in the matter of financial altruism, so that in another sense of 'luck' I am lucky to live in this society. This is what Nagel calls *luck in the problems which have to be faced*.<sup>14</sup> Still using the same example, these days I may not necessarily be praised for my contribution to Oxfam if my motivation is naively to 'feed the poor'. The debate has got more difficult than that: to the extent that 'humanitarian' assistance reinforces patronizing images of third-world dependency, or only serves the 'bureaucracy of the aid system',<sup>15</sup> then perhaps my donation is not so praiseworthy after all. Increasing awareness of 'colonialist' patterns in development work over the past ten to twenty years has increased the problems which have to be faced in deciding whether and how to give support to countries of the South. Perhaps I am actually unlucky in that? If I were an old-time philanthropist, I might not have needed to concern myself with whether the recipients of my 'charity' resented me for it. At least, it would never have occurred to me that they would.

Certainly a similar kind of 'ill-luck' in the problems which have to be faced confronts doctors when they attempt to act in the patient's best interest but find patients very much more vocal than they used to be. The problems which have to be faced are no longer just clinical (if they ever were): they are ethical as well, so that the very need to consider ethical questions raises issues of moral luck. This dilemma about paternalism versus patient autonomy, and what the moral luck paradox has to say about it, will recur in chapter 4 of this volume, on informed consent. In other ways, too, modern doctors are unlucky in the range of dilemmas which have to be faced: the rise of new technologies famously calls up new medical dilemmas all the time. The chapters of this book on new reproductive technologies, research and genetics all illustrate that point.

A very difficult set of issues concern *luck in character*, similar to Williams's category of *constitutive luck*: 'the kind of person you are, where this is not a question of what you deliberately do, but of your inclinations, capacities and temperament.'<sup>16</sup> Kant assumes that the good will may shine like a jewel in the blackest of circumstances, but what if having the good will is itself partially determined by luck? This dilemma radically undermines the notion of moral agency itself;

yet in some areas of psychiatric ethics – considered in chapter 8 – it is strikingly relevant. Similarly, to the extent that character and personality have a genetic basis, the problem of luck in character occurs there, as chapter 9 will explore. More broadly: the role of medical education, and of the Hippocratic oath, has been argued to be as much about inculcating the ‘right’ kind of character traits as about teaching knowledge or skills. What the ‘right’ kind of trait is has been seen to vary over time, from the ‘reassuring’ manner of the ‘man in charge’ to the modern-day emphasis on communication skills and consultation. But suppose none of this does any good at all? Suppose good or bad character is simply determined by luck?

Joel Feinberg argues that the Kantian attempt to drive responsibility inward, to intentions, the will or character, is ill-founded: we cannot evade the operations of luck so easily. As Margaret Urban Walker puts it:

Feinberg noted that it won’t work to drive responsibility ‘inside’ to ‘acts of the will’, ‘intentions’, or ‘character’ as a way to evade the ‘luck’ of having other factors contribute to what you have done. For chance events, other people and circumstances can influence, even determine or avert, the occurrence of ‘inner’ states as well.<sup>17</sup>

The counsel of despair here would be that luck in character so radically undermines attempts to praise or blame that we cannot really have such a thing as ethics at all. The counter-argument put by Claudia Card is worth considering, however: character is shaped by how and for what we are held responsible.<sup>18</sup> We need not simply consign character to the unexplained and inexplicable realm of ‘luck’, to which ethical judgements cannot apply. Rather, ethical judgements also shape the development of character, in a reciprocal process.

### *Outcome luck: further considerations*

The original sense in which moral luck was used was in relation to outcome luck, particularly by Bernard Williams. A recurring theme in this book is that the extent of outcome luck has been underestimated by philosophers, and that, although it occurs constantly in clinical practice, medical ethicists (and the clinicians whom they hope to assist) do not always recognize the paradox. When they do, in my experience of clinical ethics committees, they find the paradox helpful rather than irritating.

Ethicists have often failed to consider all possible outcomes, or have even assumed that the desired outcome is inevitable. Take, for example, the old saw about whether it is ethically justifiable to kill one person in order to save one hundred. The presumption is always that the hundred people will indeed be saved by the death of the one; but suppose they aren't? Even Williams himself can be accused of underestimating the number of possible outcomes in his well-known hypothetical case of Pedro and Jim.<sup>19</sup>

In this example, a traveller, Jim, happens into a South American market square where a captain, Pedro, is about to execute twenty Indians. Pedro offers Jim a chance to save all the Indians but one, whom Jim must kill by his own hand. If Jim refuses, Pedro says that all the executions will proceed; if he assents, Pedro promises that the other nineteen Indians will go free. Williams concludes rather grudgingly that the utilitarian answer – Jim should kill the Indian – is right in this case, although not for the conventional utilitarian reasons. As he puts it: 'The utilitarian is probably right in this case [but] that is not found out just by asking the utilitarian's questions.'<sup>20</sup>

Why Williams comes down in favour of Jim's killing the Indian is not the concern of my argument; I merely wish to point out that he has not fully considered what happens if things go wrong. It is all too easy – the stuff of B-grade films – to imagine Jim turning trembling to Pedro, smoking revolver in hand, and the camera panning onto Pedro's delighted grin as he takes the gun and shoots the first of the nineteen other Indians, whom he meant to kill all along.

Does this 'unlucky' outcome make Jim's decision wrong? A preliminary question that must be resolved is whether to call this outcome a mere matter of luck. First, 'unlucky' is too bland a term to describe the disaster that ensues. The more problematic matter, however, is whether luck means 'unforeseen contingency'. If so, then Jim, and Williams, cannot be castigated for failing to take it into account, for underestimating the number of positive outcomes. Here we can gain some purchase from Dworkin's distinction between choice and chance, or between option luck and brute luck. It seems inappropriate to claim that if all the Indians die, although Jim has decided to kill one of them in an attempt to save the others, it is merely a matter of brute luck or the operations of nature. Jim had a choice: this is a question of option luck. In reviewing all the options before making the choice, Jim, and Williams, failed to foresee a crucial possibility. In other words, I (and Dworkin) do not define luck as unforeseen contingency, merely as contingency.

A Kantian would assert that Jim's decision was wrong all along, regardless of the 'unlucky' consequences. More importantly, a Kantian *could* assert that the choice's rightness or wrongness is independent of the outcome. An ethical consequentialist, such as a utilitarian, would by definition be unable to do so.

Now possibly some utilitarians would argue that no additional diminution of welfare results from Jim's decision: all twenty Indians would have died in any case. But having failed to consider all the consequences should be a cause of self-mortification for any utilitarian strategy that incorporates probability assessments. This would include expected-value strategy, a common decision-theoretical model which advises the agent to choose among courses of action by multiplying the probability that an action will produce a desired outcome times the utility of that action, and its associated models, such as utilitarian economics, rational-choice theory in political science, cost-benefit analysis and quality-adjusted life years in health care economics. If the total number of possible consequences is not identified correctly, none of the probability estimates attached to any particular outcome can be valid. And this particular ill outcome should give pause to any utilitarian analysis that sets a negative value on giving additional pleasure to someone like Pedro. In that sense, Jim does make something worse happen by shooting the Indian. Certainly this unaccounted-for outcome should trouble Williams, since he sets a non-utilitarian, independent value on Jim's integrity, which has now been shattered for nothing.

In fact there are four possible outcomes in the Jim and Pedro example, not two:

- 1 Jim refuses to shoot the single Indian, and Pedro executes all twenty.
- 2 Jim refuses to shoot the single Indian, but Pedro, impressed by Jim's steadfastness and embarrassed by his own bloodthirstiness, repents and spares the others.
- 3 Jim shoots the single Indian, and Pedro liberates the nineteen others.
- 4 Jim shoots the single Indian, but Pedro breaks his promise. All twenty Indians die, as in scenario 1.

There are of course other possible outcomes, such as the cavalry rolling up in the nick of time, but these really are extraneous to the central agents, Jim and Pedro. Here I follow Williams's reminder to

consider Pedro as an agent as well: 'While the deaths, and the killing, may be the outcome of Jim's refusal, it is misleading . . . to leave Pedro out of the picture in his essential role of one who has intentions and projects, projects for realizing which Jim's refusal would leave an opportunity.'<sup>21</sup> It may be questioned, however, whether Williams obeys his own injunction. By assuming that Pedro is incapable of options 2 and 4 or, more properly, by not considering Pedro sufficiently as an agent to see that options 2 and 4 are possible, Williams does in fact leave Pedro out of the equation, the terms of which he simplifies excessively. If Pedro is seen as an agent, whose mind Jim cannot fully know, then the problem of moral luck becomes all the more complex, insofar as the outcomes for which we may be responsible hinge not only on our own actions, or the actions of an impersonal nature, but also on the actions of other agents.

In later work Williams has dismissed the possibility of outcomes 2 and 4, regarding 1 and 3 as certain beyond reasonable doubt. But outcome 2 has occurred in several sieges involving the taking of hostages, when terrorists release their prisoners even though the authorities refuse to concede. It seems that Williams's impatience is directed at a more general political argument which suggests that 'the efficacy of the detestable action is more doubtful than the example supposes.'<sup>22</sup>

This is a line often taken by those defending an absolutist position in cases of detestable actions extorted by threats made by hijackers and so forth, to the effect that the very character of the threat shows that one has reason to doubt the efficacy of giving in to it. Why should one expect such threateners to keep their promises anyway? As a *general* line of argument, this seems, bluntly, a cop-out.<sup>23</sup>

To me it seems that the 'cop-out' lies in attempting to write off possible outcomes which the moral agent ought to consider, assuming that she is considering consequences at all. However, at this stage of the book I do not want to commit myself to a general absolutist line such as Williams castigates. It may be that our disagreement is only over whether the Jim and Pedro example is one of those 'cases in which it is a reasonable bet that nothing is to be gained by giving in to threats.'<sup>24</sup> Even more narrowly, at this stage I merely want to suggest that it ought to be permissible to consider whether the Jim and Pedro example is such a case.

Suppose that outcome 2 does occur: Jim refuses to compromise his principles, and Pedro relents. The common response would be to

admire Jim and to rejoice in the happy outcome. This casts doubt on whether Jim should abandon his integrity, but is the common response correct?

Intuition allows the actual outcome to make a difference to how we assess the agent's choice, but intuition also holds that moral choice is somehow impervious to how things turn out. We commonly do admire steadfastness, but would we call Jim mule-headed or callous if Pedro did not back down? This possibility may perhaps lie behind Williams's impatience with those who counsel against giving in to threats. Yet how can the character of the agent depend on the outcome? Character is meant to be more lasting than the rush and tumble of outcomes; we generally want to think that we are more than the sum of all the outcomes that occur in our lives. Thus the common response, that Jim should be admired for steadfastness if Pedro backs down, is incoherent. It exemplifies the paradox of moral luck, and we have already seen that the paradox of moral luck threatens the very notion of responsible agency. Yet the virtue of steadfastness, which would commonly be admired if Jim remained firm, also implies that there is a wholeness to the agent that cannot be diminished by ill-luck in outcomes.

*Moral luck: how serious and genuine is the paradox?*

The impetus in the debate between Williams and Nagel over moral luck is this inconsistency in common sense, and the broader, deeply disturbing question of whether outcome luck can threaten the concept of responsible moral agency. Both begin from Kant, as the source of the view that there can be a quintessential form of value, moral value, which is 'unconditioned'<sup>25</sup> – that is, free from external contingency. Both are also concerned to distinguish outcome luck from luck in character, antecedent circumstances, or problems which have to be faced. Their motive is to limit the potentially illimitable, although in this book I will be extending the operations of risk and luck beyond luck in outcomes. This is itself a risky strategy, although less so in a book-length volume. Listing all these ways in which moral agency appears to be subject to chance leads to a heightened sense of the possible threats to the very notion of moral agency from risk and luck.

If one cannot be responsible for the consequences of one's acts due to factors beyond one's control, or for antecedents of one's acts that are

## 12 *Ethics versus Luck?*

---

properties of temperament not subject to one's will, or for the circumstances that pose one's moral choices, then how can one be responsible even for the stripped-down acts of the will itself, if they are the product of antecedent circumstances outside of the will's control?<sup>26</sup>

We have already moved a long way from Kant's assertion that the good will has full worth in itself, even if it is totally powerless to achieve the good end. What is deeply troubling about the problem of moral luck is that it undermines the notion of the moral will itself, making it impossible for us to operate as ethical agents. Even confined to outcome luck, the paradox of moral luck is deeply troubling. As Nagel puts it:

The inclusion of consequences in the conception of what we have done is an acknowledgement that we are parts of the world, but the paradoxical character of moral luck which emerges from this acknowledgement shows that we are unable to operate with such a view, for it leaves us with no one to be.<sup>27</sup>

If this is a genuine paradox, Williams agrees, then things really are quite serious, and there are only two possible resolutions, both unpalatable.

One's history as an agent is a web in which anything that is the product of the will is surrounded and held up and partly formed by things that are not, in such a way that reflection can only go in one of two directions: either in the direction of saying that responsible agency is a fairly superficial concept, which has a limited use in harmonizing what happens, or else that it is not a superficial concept, but that it cannot ultimately be purified.<sup>28</sup>

Put another way, what is serious for medical ethics about moral luck is that it radically undermines the foundational, Kantian notion of autonomy.<sup>29</sup> Critiques of extreme advocates of personal autonomy have charged the autonomy view with a failure to count relationships in, to include 'the world' (in Nagel's terms) or 'the web' (in Williams's words).<sup>30</sup> But the opposite extreme apparently leaves no room for moral agency or personal autonomy at all; indeed, it seems to make the entire debate a nonsense, since there can be no such thing as ethics, medical ethics included, without a notion of responsibility and moral agency. If we are ultimately responsible for everything that happens in the world, then we are responsible for nothing.

But are matters really this serious? Is moral luck a genuine contradiction in terms, a true paradox, or has Williams overstated the problem? As a preliminary way of examining Nagel's answer to Williams, let us recast Williams's formulation of the moral luck dilemma in other terms, borrowing from Michael Zimmerman's method of restating the puzzle.<sup>31</sup>

- 1 A person *P* is morally responsible for an event *e*'s occurrence only if *e*'s occurring was not a matter of luck.
- 2 No event is such that its occurrence is not a matter of luck.
- 3 Therefore, no event is such that *P* is morally responsible for its occurrence.

Nagel, according to Zimmerman, denies the conclusion in (3) but accepts both premises (1) and (2), believing the paradox to be genuine (as indeed he must, to accept the first two premises but deny the resultant conclusion). Williams tends to accept the second premise while denying the first. If this is so, then Williams is not obliged to accept the paradox as genuine – still less to call for a complete overthrow of the Kantian ethical tradition, as he does, because he believes it is radically undermined by the paradox. In fact Williams goes so far as to doubt whether any doctrine of ethics, Kantian or non-Kantian, can resolve the paradox, although he holds out hope for Aristotelianism.<sup>32</sup> I suspect that Williams does accept the first premise as well. Certainly he is no less sceptical than Nagel, as Zimmerman's assessment would imply. In fact Williams presents himself as being more sceptical than Nagel about the possibility of an unconditioned form of moral value, untainted by luck.<sup>33</sup>

Both Nagel and Williams, however, appear to accept premise (2), which I think is actually the weakest link. *Although all events are a matter of luck, they are not all a matter of moral luck, nor are they all moral events.* Much of this book will be dedicated to examining this 'escape route' in relation to practical problems in medical ethics.

Zimmerman also proposes that we narrow the terms down; he believes that recasting the terms dispels the moral luck problem. Thus, if premise (1) is reformulated as '*P* is morally responsible only if she or he was in restricted control of *e*', premise (2) cannot be upheld – if it is in turn recast as 'No event is such that anyone is ever in restricted control of it.' To reverse the procedure, premise (1) becomes palpably false if refashioned as '*P* is morally responsible only if she or he was in unrestricted control of *e*,' while (2) is true if

revised to read ‘No event is such that anyone is ever in unrestricted control of it.’ Thus the distinction comes to hinge on restriction over control.

In medical ethics, this recasting has a certain appeal. It builds in the rather unique way in which doctors’ control is restricted not only by brute luck – the nature of illness – or by option luck – their own choices – but also by the choices of others – members of the clinical team, health care managers, patients and families. (In Dworkin’s typology, brute luck probably includes the choices of others, but this manoeuvre does not really do justice to the decisions of others as moral agents.) Restriction over control is the everyday stuff of medical decisions. Yet it is palpably false, at least in law, to say that ‘Doctor *D* is not responsible for failing to obtain informed consent, and for the ensuing battery to patient *P*, because she was not in unrestricted control over the course of events.’ Much tension in medical ethics and law arises from the interplay of restriction and control, for example, in medical negligence, resource allocation, and research ethics. But even if the paradox is not so crippling as Williams believes, adding the refinement about restricted control does not dispel it altogether, I think.

To resolve the paradox of moral luck, we could absolve agents from responsibility for all matters beyond their ability to predict or control. But then we will wind up holding them responsible for very little. This is the strategy which Williams has pursued in subsequent work.<sup>34</sup> On a purely practical level, this tactic seems ill-advised in medical ethics: it leaves patients with very little protection and doctors with very little responsibility. On a considerably deeper level, it also denies the moral agency of doctors and makes them less complete people. On an intermediate level, it leaves doctors with no other guide than hindsight, which is of course no guide at all.<sup>35</sup> To understand why this is so, we need to examine Williams’s examples, which illustrate his assertion that only success can justify an agent’s decision under uncertainty.

### *Judgement from hindsight: Gauguin and Anna Karenina*

Williams’s first example is a fictionalized version of the painter Gauguin, who abnegates his responsibility to his wife and children when he abandons them to paint in the South Seas. Gauguin’s failure as a painter would prove his decision to abandon his family was wrong in this rather opaque sense:

If he fails . . . 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.<sup>36</sup>

What can Williams mean here? Far from being so obvious as to be platitudinous, the rightness of Gauguin's action if he succeeds would seem quite hard to defend. If only success can justify a decision, but success is not certain at the time the agent makes a choice, there will turn out to be no basis but hindsight for judging whether an action was morally right or wrong. We do not generally think that agents are responsible for matter that they could only have known with hindsight. Therefore, if we take this view, we will not be able to hold agents responsible for very many of their choices.

Williams recognizes that Gauguin's success or failure cannot be predicted in advance; what he seems to be claiming is that only success can show us whether Gauguin had the sort of talent that could have served as a basis for the thought, at the time the decision had to be made, that his talent was the genuine article. In addition, he distinguishes between the factors intrinsic to Gauguin's success or failure and those extrinsic to it. Gauguin's decision is subject to luck on both counts, but only the intrinsic factors count in proving his decision to have been wrong.

From the perspective of consequences, the goods or benefits for the sake of which Gauguin's choice was made either materialize in some degree, or do not materialize. But it matters considerably . . . in what way the project fails, if it fails. 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.<sup>37</sup>

Williams is distinguishing between the failure of Gauguin's project and his failure as a moral agent. If things go wrong, they can go wrong in two ways, similar to the ways called option luck and brute luck by Dworkin. In Williams's view, Gauguin is not morally

blameworthy if he abandons his family but fails to achieve his goal for reasons of brute luck. An event such as injury 'is too external... to unjustify him, something which only his failure as a painter can do.'<sup>38</sup> These external or extrinsic causes cannot always be predicted, but Gauguin could and should have sounded his ambitions and talents sufficiently to know whether he had it in him to be a great painter. If he failed because he did not have those qualities, then he could have been blamed. However, he may also have bad intrinsic luck, which is much more serious. The Gauguin who must assess his own talent as an artist, back in Paris, is one who has not yet sounded the full range of his artistic talents; he is in an important sense a different person from the Gauguin he might become if he goes to Tahiti. Since his knowledge is imperfect, and since he can know only in hindsight, as the Tahiti Gauguin, whether his decision as the Paris Gauguin was justified, the problem of moral luck in the decision he must make in Paris is serious and inescapable.

In chapter 6 I draw on a similar dichotomy in discussing how medical professionals should feel if they choose one patient over another for receipt of a scarce medical resource on the basis of the first patient's better prognosis, but if the patient who receives the kidney or the expensive procedure then dies. The doctors chose the 'wrong' patient, but it is inherent in the nature of prognoses that a certain proportion of them will turn out 'wrong', since they are statistical judgements. I present this example as a moral luck case in chapter 6, but I also suggest principles that would allow it to count as a project failure rather than a personal failure for the doctors as moral agents. Thus I do accept and draw on elements of Williams's analysis, but ultimately I agree with Nagel that it is too crippling.

Not all kinds of ill-luck are moral. Williams relies heavily on the idea that, if Gauguin does turn out to be a great artist, that outcome carries moral weight. But what if art has nothing to do with morality? In that case Gauguin's moral decision – and abandoning his family does incontrovertibly seem a moral choice – cannot be justified or unjustified on the basis of his artistic success or failure. Moral luck cannot enter into it: the decision is right or wrong from the start, regardless of outcomes, and no hindsight is required. Nagel's reply to Williams runs along similar lines:

[Williams] points out that though success or failure cannot be predicted in advance, Gauguin's most basic retrospective feelings about the decision will be determined by the development of his talent. My disagree-

ment with Williams is that his account fails to explain why such retrospective attitudes can be called moral. If success does not permit Gauguin to justify himself to others, but still determines his most basic feelings, that only shows that his most basic feelings need not be moral. It does not show that morality is subject to luck.<sup>39</sup>

Gauguin's situation may be tragic if he cannot fulfil his genuine artistic talent without doing something morally wrong. But that is not the same as a genuine moral dilemma, in which both courses are morally wrong: for example, Sophie's choice, in William Styron's novel of that name, between allowing both her son and her daughter to die at the hands of the Nazis and choosing which one of them will live. Such moral dilemmas may also be subject to moral luck – Sophie cannot know in advance that the son, whom she has chosen, will die of natural causes in the concentration camps – or they may in fact be a subset of moral luck dilemmas.<sup>40</sup> But either way Gauguin's choice is not a moral decision, and therefore, by definition, it is not vitiated by moral luck. If he succeeds, he cannot justify himself to his abandoned wife and family by his success; if he fails, he is not more to blame because he has not succeeded – he was already to blame.

A similar criticism can be levelled at Williams's second example, that of Anna Karenina: it is not specifically about *moral* luck. Williams adds the (similarly schematized) example of Anna Karenina to that of Gauguin in order to illustrate the more typical case of someone for whom the locus of intrinsic luck is partly external. Here we begin to see important ways in which Williams's typology, intrinsic and extrinsic, differs from Dworkin's option and brute luck.

The intrinsic luck in Gauguin's case concentrates itself on virtually the one question of whether he is a genuinely gifted painter who can succeed in doing genuinely valuable work. Not all the conditions of the project's coming off lie in him, obviously, since others' actions and refrainings provide many necessary conditions of its coming off – and that is an important locus of extrinsic luck. But the conditions of its coming off which are relevant to unjustification, the locus of intrinsic luck, largely lie in him – which is not to say, of course, that they depend on his will, though some may. This rough coincidence of two distinctions is a feature of this case. But in others, the locus of intrinsic luck... may lie partly outside the agent, and that is an important, and indeed the more typical case.<sup>41</sup>

Williams takes Anna Karenina's suicide to be an admission that she has failed, not merely that her project has failed through some extrinsic factor – for example, if Vronsky had been killed in an accident. 'What she did, she now finds insupportable, because she could have been justified only by the life she hoped for, and those hopes were not just negated, but refuted, by what happened.'<sup>42</sup>

If we accept Williams's interpretation of the Anna Karenina story, grave paradoxes certainly ensue. Anna could not have known in advance that her affair would end badly, but she could only have been justified by its turning out well. What basis could she then have had for her decision about whether to leave her husband? This looks at first to be the problem about hindsight and responsibility that I identified before.

However, this example is really about what counts as an ethical choice and what agents are responsible *for*. The case of Anna Karenina cannot be the universal example that Williams thinks it is, because in Anna's society only women had to make such a choice. (Admittedly Anna is a schematized figure, as was Gauguin, but this excuse does not get round the problem of non-universalizability.) There is no moral luck in this example, because there is no genuine moral choice or responsibility.

Here the Kantian position would suggest that the genuine moral life must, by its nature, be equally open to all and its dictates be equally applicable to all. Because men were free of the strictures that required Anna to make a choice, because they operated under a much less restrictive system of responsibilities than did women in Anna's society, that system cannot be accounted a genuine moral system; nor can Anna's failure to live up to it count as a moral failure. Indeed, the fact that she had to make a decision at all is not evidence of a moral choice, but of victimization.

Anna had to decide between her lover and her husband and son. Her brother Oblonsky was allowed to keep both his mistresses and his home life. As the novel quite pointedly tells us at the very opening, 'although Oblonsky was entirely in the wrong as regards his wife, as he himself admitted, almost everyone in the house, even the nurse, [his wife's] best friend, was on his side.'<sup>43</sup> With a fine and deliberate irony, the novel begins with Anna persuading her sister-in-law to forgive Oblonsky for his latest dalliance. Williams seems to want to argue that Anna's affair was wrong from the start because it ended badly. Her brother's affairs don't end badly: are they any less wrong?

Williams does hone his analysis by noting that, once Anna left her husband, the relationship with Vronsky would have to bear too much weight. He calls her failure to foresee this an intrinsic matter. Certainly the novel's description of the lovers' scenes confirms that the relationship is overloaded, but that is only because society will not allow Anna to retain the marital relationship as well.<sup>44</sup> Although Williams calls this failure to forecast the extra burden on the relationship 'a truth not only about society but about her and Vronsky',<sup>45</sup> I cannot see that the occasion of the so-called moral decision would have arisen but for society's extrinsic laws and *mores*.<sup>46</sup> Still less do I agree with Judith Andre's contention that we would be right to praise Gauguin and blame Anna Karenina because 'the person who can correctly assess his or her chances of success is better formed than the person who cannot.'<sup>47</sup> This is to confuse acting ethically with being a clever technician. In any case, we do not know whether Gauguin judged the probability of success accurately or inaccurately: we only know that he succeeded.

### *Escaping from the paradox*

The examples of both Anna Karenina and Gauguin can be construed as stories about something other than *moral* luck, and thus as no real threat. In most of this book I will be arguing along similar lines: that we can retain a notion of responsible agency, one which is not ultimately unsettled by moral luck, if we limit what agents are responsible *for*. Thus in chapter 4 on informed consent to treatment, I want to say that doctors are not responsible for the bad outcome of a procedure, provided they are not negligent, but that they are responsible for the prior duty of ensuring that the patient is informed of the likelihood of the procedure turning out wrong. In this way, and this way alone, responsibility for ill-luck in outcomes is transferred from the doctor to the patient. Although fully informed consent is something of a chimera, and a heavy duty in itself, obtaining as genuine and mutually communicative a consent as possible in the circumstances lessens the doctor's duties in the long run. Without such a notion, doctors would blame themselves for each procedure that turns out badly; no one can practise for long with such a mindset.

There is another possible escape route from the paradox which may strike a chord with many doctors, immersed as they are in

history-taking: a narrative approach to the moral luck dilemma. This is what Margaret Urban Walker means, I think, when she writes:

The agent is not a self-sufficient rational will fully expressed in each episode of choice, but is a history of choices . . . for whom episodes are meaningful in terms of rather larger stretches . . . We ought not to be surprised that . . . pivotal episodes which give sense to large segments are adequately judgeable only in retrospect.<sup>48</sup>

Chapter 9 will consider a developmental approach to responsibility, similar to Walker's approach. We will see there that such an approach is particularly relevant to medical genetics.

I have criticized Williams's two examples, but I firmly believe that medical ethics is full of genuine moral luck cases. It should be possible to stipulate in advance that good examples must involve agents who stake their moral success – not their worth as artists (or clinicians), their happiness in affairs foredoomed by social discrimination, or any other value – in advance on outcome luck.

Here is another such example: a research ethics committee that approves random clinical trials of a new drug on bowel cancer patients without informed consent. A patient with a prognosis of several years of active life dies two weeks later, because her bone marrow becomes irreversibly depressed in function after administration of the drug.<sup>49</sup> Let us suppose that no medical negligence is involved, that the drug was administered correctly but turned out to have side effects that no one could foresee. Is the outcome of the experiment the committee's fault? Or are there ways in which research ethics committees can limit what they are ultimately responsible for, and so avoid the paradox of moral luck?

It is important for doctors not to allow luck, risk and uncertainty in ethical decision-making to weigh so heavily on their minds that they are ultimately unable to act at all.<sup>50</sup> This is itself a risk, one that Williams rather too willingly embraces. Perhaps because I am not convinced about Williams's two examples, I am not persuaded by his fear of the 'final destruction' that occurs when Kantian strictness about the purity of the will is united with a utilitarian doctrine of negative responsibility. 'There is, at the end of that, no life of one's own, except perhaps for some small area, hygienically allotted, of meaningless privacy. Because that is a genuine pathology of the moral life, the limitation of the moral is itself something morally important.'<sup>51</sup>

Here I do agree with Williams, but I believe that he is now pointing towards a different path out of the maze. That is, rather than saying that we can have no meaningful conception of moral agency, he is claiming that we can have such a notion if we limit what agents are responsible for. This is close to my own hypothesis. An extensive catalogue of kinds of *luck* – such as the four suggested earlier in this chapter – is less crippling to the notion of moral agency than is a very broad notion of what counts as *moral*. No matter how deeply Gauguin feels about his failure as an artist, if he should fail, the intensity of that feeling does not itself make the failure a matter of ethics. If it is an ethical decision, that is because Gauguin breached his responsibilities to his family. They may or may not forgive him; they may cease to blame him, whether or not he succeeds as an artist. But if he stops being blameworthy, it is not because he succeeded as an artist; his original decision to abandon them remains a breach of responsibility.

Is this merely juggling with terms? I think not, in the example of Gauguin; but I also think that there will be cases in which the paradox of moral luck is genuine and troubling. Negligence is just such an area. Nagel gives the example of a lorry driver whose careless driving kills a child who darts into the road at the wrong moment. If the child had not run into the road, she would not have been killed; but, equally, the driver's negligence deserves blame. The question is how much blame. Here the law is somewhat equivocal, as is intuition. Many of us would blame the driver more if the child is killed than if she is not; yet this is to allow the operations of outcome luck to determine our moral judgements. If the driver is negligent, he is negligent not in the consequences but in his failure to drive properly, for which he can be blamed regardless of whether or not a child dies.

Perhaps we, and the law, do increasingly take this line in relation to drink-driving, but it does seem somewhat unrealistic to demand that anyone who exceeds the safe drinking limit should feel as deeply and profoundly to blame if he does not kill someone as if he does. Similarly, is the overworked junior doctor who prescribes the wrong dosage after an eighty-hour work week equally to blame if an experienced nurse pulls him up for his error, or if no one intervenes and a patient dies? It is unlikely that he would be disciplined by the General Medical Council in the first case, but should he be? If he and all his ilk were, would we eventually have no doctors at all?

On the other hand, it is surely correct to say that the drink-driver's negligence lies in driving while drunk, whether or not he causes an accident. He is being negligent in subjecting others' lives to unnecessary risk,<sup>52</sup> but what do we mean by 'unnecessary'? Exactly this question arises in many areas of medical ethics, particularly in research and psychiatry, and I will deal with it at greater length in chapter 8, as well as in chapter 4, which examines rationality and risk assessment in the context of consent to treatment. Virtually all our actions subject someone to risk, just as everything can be seen to hinge on luck, but in both cases these formulations are simultaneously insightful and too broad. To assert that we should never subject anyone to any added risk seems implausible, particularly in medical research; yet to use that as an excuse for piling on additional risks also seems an untrustworthy argument.

A firm critic of the notion of moral luck would have to assert that we should treat both sorts of case equally. How then can we explain the difference we draw between them? – except to say that we are wrong to do so. Norvin Richards, for example, claims that the ill-luck of the lorry driver who does kill the child essentially lies in getting caught. The driver who kills the child makes it clear to the world that he is a negligent driver; he should actually count himself lucky that he was not disciplined beforehand. Thus, luck does not affect one's responsibilities or deserts, only our knowledge of them.<sup>53</sup> Whereas Nagel accepts that moral luck is a genuine paradox (even if it does not apply to the wide range of events Williams wishes it to) and that the negligent driver who kills is more to blame than the one who does not, Richards asserts that the paradox is not genuine but merely a matter of confused thinking. Luck is more a matter of our epistemic position, our uncertainty and ignorance, than of what actually determines an agent's deserts. We must allow an irresponsible agent's luck in getting caught or not to determine to some extent how we treat him, since we are not omniscient. But that does not mean that ethics is radically undermined by luck in the way that Williams and Nagel claim it is.

This particular escape from the conundrum strikes me as rather cynical, although somewhat attractive, in that it logically calls on the practitioner to take seriously those examples of bad practice which are not detected, if only from self-interest. More deeply, it may be attractive insofar as it affords us a way out of the conclusion reached by both Nagel and Williams: that moral luck is a genuine paradox, imposing varying degrees of handicap on our ordinary assumptions

about morality. A third possibility should also be considered: that moral luck is real – not merely a matter of fuzzy concepts – but that it is not a paradox. Rather, incorporating moral luck into our everyday assessments is not only possible but desirable. This is the virtue-centred view taken by Margaret Urban Walker,<sup>54</sup> which, together with similar proposals by Martha Nussbaum and Alasdair MacIntyre, will be examined at greater length in chapter 2.