

Moral Luck: A Partial Map

Author(s): Michael J. Zimmerman

Source: *Canadian Journal of Philosophy*, Dec., 2006, Vol. 36, No. 4 (Dec., 2006), pp. 585-608

Published by: Cambridge University Press

Stable URL: <https://www.jstor.org/stable/40232303>

---

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact [support@jstor.org](mailto:support@jstor.org).

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <https://about.jstor.org/terms>



Cambridge University Press is collaborating with JSTOR to digitize, preserve and extend access to *Canadian Journal of Philosophy*

JSTOR

## *Moral Luck: A Partial Map*

MICHAEL J. ZIMMERMAN  
University of North Carolina at Greensboro  
Greensboro, NC 27402-6170  
USA

Luck varies from person to person, for two reasons. First, for something to occur as a matter of luck is for it to occur beyond the control of someone, and what is beyond one person's control may not be beyond another's. Second, luck may be either good or bad (or neutral — but in that case it is not very interesting), and what is good luck for one person may be bad luck for another.

Moral philosophers have paid a good deal of attention to luck in an effort to determine its relevance to moral judgments of various sorts. We may distinguish three broad classes of such judgments: *aretaic* judgments, having to do with moral virtue and vice; *deontic* judgments, having to do with moral obligation; and what I will call *hypological* judgments,<sup>1</sup> having to do with moral responsibility. In the wake of Harry Frankfurt's ground-breaking discussion of the claim that moral responsibility requires control,<sup>2</sup> most of the attention that has been devoted in recent years to moral luck has concerned hypological judgments in particular; deontic judgments and aretaic judgments have been given relatively short shrift. In this paper I will attempt a partial rectification of this bias; although I will not concern myself further with aretaic judgments, I will have a good deal to say about the relation between moral obligation and luck.

The plan of the paper is this. Section I: I seek to clarify the distinction between hypological and deontic judgments. The failure fully to appre-

---

1 The term is drawn from the Greek *ὕπολογος*, meaning 'held accountable or liable.'

2 See Harry G. Frankfurt, 'Alternate Possibilities and Moral Responsibility,' *Journal of Philosophy* 66 (1969) 829-39.

ciate this distinction explains in part, I think, the tendency in recent discussions of moral luck not to attend explicitly to the relevance of luck to deontic judgments. Section II: I draw five distinctions concerning how something may be within or beyond someone's control, three of which I select for further examination. Section III: I consider the question whether moral obligation and moral responsibility differ with respect to how they relate to luck in terms of the first of these three distinctions: my tentative answer is 'No.' Section IV: I consider the question whether moral obligation and moral responsibility differ with respect to how they relate to luck in terms of the second of these three distinctions: my answer is 'Yes.' Section V: I consider the question whether moral obligation and moral responsibility differ with respect to how they relate to luck in terms of the third of these three distinctions: my answer is again 'Yes.' Section VI: I consider the relevance of luck to moral dilemmas. Section VII: I return to the second of the three distinctions and note why it is that, in light of the discussion in Section VI, the difference between moral obligation and moral responsibility with respect to how they relate to luck in terms of this distinction is very deep — deeper than acknowledged in Section IV. Section VIII: I provide a summary of my findings.

As the foregoing plan indicates, what follows will be more in the nature of a survey than of a detailed exploration. This is because there is a lot of ground to cover, much of it hitherto wholly uncharted. The map that will emerge is therefore not only partial (since it entirely overlooks aretaic judgments) but large-scale. This may be disappointing, in that many important questions will not receive the close attention that they deserve. Nonetheless, I take my project in this paper to be a necessary preliminary to bringing these questions into sharper focus. My hope is that my findings will prompt and guide further investigation into the relevance of luck to moral judgments of all sorts.

## I

It seems often to have been assumed that what is to be said about the relevance of luck to hypological judgments can be carried over wholesale to what is to be said about the relevance of luck to deontic judgments. This may be due to a second assumption, according to which one is morally responsible for having done something if and only if one had a moral obligation not to do that thing but did it nonetheless. This claim may seem especially tempting if one puts it entirely in the terminology of responsibility: one is (retrospectively) responsible for having done something if and only if one was (prospectively) responsible for not doing it but did it nonetheless. Such a dual use of 'responsible' is liable to mislead. Henceforth, I will use the term only to refer to what is often

called retrospective responsibility. I will not talk again of prospective responsibility; instead, I will talk of obligation. The second assumption should be rejected, though, even when it is restricted to overall moral obligation (which is my topic in this paper) and thus excludes merely *prima facie* obligation.

One problem with the second assumption is that it confuses responsibility in general with culpability in particular; it overlooks laudability, the 'positive' aspect of responsibility. Whereas it may seem plausible to say that one is culpable for having done something if and only if one had an obligation not to do it that one did not fulfil, it is surely not at all tempting to tie laudability in this way to unfulfilled obligations.<sup>3</sup>

Another problem is that, even when it is restricted to culpability, the second assumption is false — in both directions, I believe. It is false to say that one is culpable for having done something *if* one had an obligation not to do it that one did not fulfil, because it is possible to have an excuse for wrongdoing. This is commonly, though not universally, acknowledged. And it is, I think, false to say that one is culpable for having done something *only if* one had an obligation not to do it that one did not fulfil, because it is possible to have what I have elsewhere called an 'accuse' despite having done no wrong. This is not commonly acknowledged, although a few philosophers have recognized the possibility.<sup>4</sup>

The reason why excuses and accuses are possible is that whether one is responsible for having done something is a function (in part) of what one *believed* one's obligations to be rather than of what one's obligations *in fact* were. One has an excuse, despite having done wrong, if one did not believe that one was doing wrong (and one is not culpable for this failure of belief). One has an accuse, despite having done no wrong, if one believed that one was nonetheless doing wrong (and one satisfies

---

3 It would be more plausible to tie laudability to fulfilled rather than unfulfilled obligations, but if there is such a link (which I doubt, for reasons analogous to those concerning the possibility of 'accuses,' regarding which see the remarks that immediately follow), it is not at all straightforward. First, it seems clear that one can fulfil an obligation in such a way that one is not laudable — i.e., one does not deserve any praise — for doing so. Second, in some contexts laudability has essentially to do with supererogation — going beyond one's obligation — rather than simply fulfilling an obligation.

4 See, for example, G.E. Moore, *Ethics* (Oxford: Oxford University Press 1965), 82; Derek Parfit, *Reasons and Persons* (Oxford: Clarendon Press 1984), 25; Judith Jarvis Thomson, 'Self-Defense,' *Philosophy and Public Affairs* 20 (1991), 295; Ishtiyaque Haji, *Appraisability: Puzzles, Proposals, and Perplexities* (Oxford: Oxford University Press 1998), 146.

certain other necessary conditions of responsibility). Responsibility and obligation are in this way wholly independent of one another, not conceptually (since responsibility is in part a function of belief about obligation) but logically (since one can be responsible in the absence of having been obligated, and *vice versa*).

There is thus a fundamental disparity between hypological and deontic judgments.<sup>5</sup> This disparity is sometimes said to reside in the fact that judgments about obligation consist in evaluations of *acts* whereas judgments about responsibility consist in evaluations of *agents*. This claim is rather misleading, however, for at least two reasons. First, both kinds of judgment — deontic and hypological — concern both agents and acts; after all, it is *agents* that are obligated to act in certain ways, and it is for their *acts* that agents may be held responsible. Second, both kinds of judgment concern not just acting in particular but behavior in general (including omission) and the consequences of behavior. Despite these shortcomings, the claim is nonetheless helpful, I think, for it directs our attention to an important truth, namely, that deontic judgments constitute what may be called *act-based* (or, perhaps, *act-focused*) evaluations, whereas hypological judgments constitute *agent-based* evaluations. By this I mean that the primary concern of deontic judgments is the assessment of agents' behavior (and the consequences of this behavior<sup>6</sup>), relative to behavior in which they could otherwise engage or have engaged. The primary concern of hypological judgments is, by contrast, the assessment of agents in light of the behavior in which they have engaged.

I have claimed that both excuses and accuses are possible and hence that hypological and deontic judgments are logically independent of one another. If I am right, there is surely good reason to suspect that what is to be said about the relevance of luck to hypological judgments cannot after all be carried over wholesale to the question of the relevance of luck to deontic judgments. I believe that this suspicion is correct, for reasons

---

5 In *Deontic Morality and Control* (Cambridge: Cambridge University Press 2002), Ishtiyaque Haji pursues in detail certain implications of the logical independence of responsibility and obligation. I endorse much, though not all, of what he says. However, he says little about the issues that I will raise in this paper (although see note 32 below).

6 Note that among these consequences may be the development (or maintenance) of certain traits of character. In this way, an agent may have an obligation to become (or remain) a certain kind of person. Even in such a case, though, the judgment that the agent has the obligation in question may be said to be 'act-based' rather than 'agent-based,' in that the focus is not on how the agent is to be evaluated but on how the agent ought to behave.

that I hope to make clear later. But I have not argued for the claim that both excuses and accuses are possible,<sup>7</sup> and you may reject it, insisting that the second assumption, when restricted to culpability, is true — that is, that one is culpable for having done something if and only if one had an obligation not to do it that one did not fulfil. But even if you do wish to insist on this, you should recognize that hypological and deontic judgments are nonetheless distinct. As indicated earlier, obligation is essentially a *prospective* matter; if one is obligated at some time *t* to perform some act at some time *t\**, then *t\** cannot be earlier than *t*. In the interval (if any) between *t* and *t\**, the obligation is and must be *unresolved*, in that whether it will be fulfilled awaits determination. If one does fulfil it, one will have done right; if not, one will have done wrong. At the moment that this issue is resolved, the obligation ceases to be. There thus arise two distinct questions concerning the relation between luck and moral obligation: (A) whether and how the *incurring* of obligation is subject to luck; (B) whether and how the *fulfilment* of obligation is subject to luck. Responsibility, by contrast, is not a prospective matter. It is and can only be *resolved*, in that the 'mode' of one's responsibility — whether that of laudability or culpability (or, perhaps, neither<sup>8</sup>) — is always determinate. There thus arises just one pertinent question concerning the relation between luck and moral responsibility: (C) whether and how the *incurring* of responsibility is subject to luck. If you subscribe to the second assumption, you will match the incurring of culpability not with the incurring of obligation but with its resolution. That is, you will hold that one incurs culpability for doing something when and only when one fails to fulfil an obligation not to do it that one earlier incurred. You will thereby be treating responsibility as essentially retrospective. If, like me, you do not subscribe to the second assumption, perhaps you do not conceive of responsibility as essentially retrospective. (I will return to this issue in Section VII.) Nonetheless, you must, I think, acknowledge that all ascriptions of responsibility are determinate, in that they constitute ascriptions of either laudability or culpability (or, perhaps, neither). It makes no sense to ascribe unresolved responsibility to someone.

In what follows, I will attend to each of questions A-C.

---

7 I have done so in Michael J. Zimmerman, 'Another Plea for Excuses,' *American Philosophical Quarterly* 41 (2004) 259-66, and Michael J. Zimmerman, 'A Plea for Accuses,' *American Philosophical Quarterly* 34 (1997) 229-43.

8 See Michael J. Zimmerman, *An Essay on Moral Responsibility* (Totowa, NJ: Rowman and Littlefield 1988), 61-2.

## II

Many philosophers claim that moral responsibility is immune to luck inasmuch as one cannot be morally responsible for something that was not in one's control in some way and to some extent. I agree. (It is important to recognize that even Frankfurt agrees. His challenge concerns the traditional understanding of the way in which and the extent to which moral responsibility requires control; he does not deny that there is such a requirement.) We might try explaining this fact by appealing to the unfairness of holding someone morally responsible for something that was not in his (or her) control.<sup>9</sup> Such an explanation might be used to argue for a disparity between responsibility and obligation as they relate to luck. It might be claimed that it is precisely in virtue of the fact that hypothetical judgments are agent-based that the possibility of unfairness in such judgments arises. Since deontic judgments are not agent-based, this possibility does not carry over to them. There would seem to be no unfairness in maintaining that someone had an obligation neither the incurring nor the fulfilment of which was in his control, as long as we do not hold him responsible either for its being incurred or for its eventual resolution.

I do not endorse this line of thinking. I grant that it is unfair to hold someone responsible for something that was not in his control in any way or to any extent, but I think this is explained by, rather than an explanation of, the fact that responsibility requires control.<sup>10</sup> I do not have a ready explanation for this fact in turn. I do not know what argument may be given for it (although I have elsewhere defended it against objections<sup>11</sup>); I will simply presuppose it here, with what I hope is some plausibility. And I think that, at least with respect to its resolution if not with respect to its being incurred, obligation is similar: no one can be obligated regarding something that is not in his control in any way or to any extent, for 'ought' (the 'ought' that expresses moral obligation) implies 'can' (in a sense of 'can' that expresses some measure of personal

---

9 This idea applies straightforwardly to judgments about culpability, but it applies to judgments about laudability too. To hold someone laudable for something he could not control is unfair, even if it is not unfair *to* that person (in that it does not constitute an adverse judgment about him).

10 Contrast R. Jay Wallace, *Responsibility and the Moral Sentiments* (Cambridge, MA: Harvard University Press 1994), ch. 4.

11 See Zimmerman, *Essay*, 113-19, and Michael J. Zimmerman, 'Taking Luck Seriously,' *Journal of Philosophy* 32 (2002), 569-70.

freedom or control). Again, I do not know what argument may be given for this claim (although I have elsewhere defended it against objections<sup>12</sup>); I will simply presuppose it here, with what I again hope is some plausibility.

Despite the fact that responsibility and obligation are alike in the way just indicated, disparities nonetheless exist between them with respect to luck, both because of the falsity (as I see it) of the second assumption noted in the last section and because of the distinction between the incurring and the resolution of obligation. (If the second assumption is true, then whatever measure of control is required in order for one to be culpable for something will be required in order for one to fail to fulfil an obligation, that is, in order for one to do wrong. If the assumption is false, then even if culpability and wrongdoing both require some measure of control, they may differ with respect to the details of their respective requirements. Furthermore, even if the assumption is true, whatever measure of control is required regarding the resolution of obligation may not be required regarding its being incurred.) In order to assess these disparities, we should distinguish certain ways in which something may be within or beyond someone's control.

First, we should distinguish *partial* from *complete* control. One has complete control over something only if its occurrence is not contingent on anything that is beyond one's control. This is obviously unattainable by anyone at any time. I mention it only to dismiss it. Any control that we ever have over anything is at best partial. (If you doubt this, consider the simple fact that whatever control you enjoy over anything depends on your having been born — something that we may hope was in someone's control, but not yours.) I assume that this fact alone does not suffice to show that we are never responsible for or obligated regarding something. The control that is required for either incurring responsibility or fulfilling an obligation is therefore merely partial.

Second, we should distinguish *basic* from *comprehensive* control.<sup>13</sup> The idea is familiar, even if the terminology is not. It can best be conveyed by means of an example. Suppose that Bert holds up a bank. He does so by pointing a gun at Alf, a teller, and ordering him to hand over all the money in his till. Alf complies, and Bert walks off with the money. Alf's manager uncharitably complains to Alf, 'What business did you have

---

12 See Michael J. Zimmerman, *The Concept of Moral Obligation* (Cambridge: Cambridge University Press 1996), ch. 3.

13 See Zimmerman, *Essay*, 24-6, in which I used instead the terms 'strict' and 'broad.'



giving all that money to that fellow? He doesn't even have an account with us!' To which Alf responds, 'He was pointing a gun at me! I had no choice!' We can all sympathize with Alf's response, but is it accurate?

Strictly, it is not. Alf did indeed have a choice — and, we may assume, he made the right one. He thus had what I call basic control over his behavior. Nonetheless, insofar as he was strongly coerced to make the choice that he made, his behavior was not fully voluntary; it was not fully free, and so he did not have what I call fully comprehensive control over it. Such control can clearly come in degrees, depending on the force of the pressures that bear on the agent.<sup>14</sup> It is often claimed that some measure of comprehensive control beyond merely basic control is required for responsibility. This may seem plausible. After all, who (other than his manager) would want to hold Alf responsible for handing over the money to Bert? But this betrays a confusion that I mentioned earlier. Perhaps it would indeed be a mistake to say that Alf is culpable for his behavior, but that is not to say that he is not responsible for it. He may in fact be laudable for it.

Regardless of whether some measure of comprehensive control is required for responsibility, when it comes to comparing obligation with responsibility it is on merely basic control that I want to focus here. This is not because comprehensive control is not relevant to judgments about obligation. On the contrary, it surely can be. (In the absence of Bert's threat, Alf is presumably obligated not to hand over the money; in the presence of the threat, his not handing it over is rendered either supererogatory or wrong.) It is simply because *luck* is a function only of basic control. (The sort of control that Alf lacks regarding his handing over the money does not make his handing it over a matter of luck.)

A third distinction is between *deliberate* and *coincidental* control. The term 'coincidental control' may seem an oxymoron, but the phenomenon to which I am applying it is genuine. Suppose that there is a safe whose contents you would dearly like to get your hands on. These contents are protected by a combination lock. To open the safe, all you have to do is turn the dial in accordance with the pertinent combination of numbers. But you do not know what the combination is. Can you open the safe? Well, yes and no. You *can* open it, since you can turn the dial in the requisite way; but you cannot open it *intentionally*. Thus you do not have deliberate control over opening it, although you do have

---

14 It may be that no one ever has fully comprehensive control over his behavior. Whether this is so depends on just what types of pressures should be thought to compromise such control.

coincidental control over doing so.<sup>15</sup> In light of your lack of deliberate control over opening the safe, it would certainly be lucky for you (and unlucky for its owner) if you managed to open it nonetheless. The relevance of such luck to both moral obligation and moral responsibility is an issue that I will address in Section III.

A fourth distinction is between *direct* and *indirect* control. One has indirect control over something just in case one has control over it by way of having control over something else. One has direct control over something just in case one has control over it that is not indirect. For example, a thoroughly drunk person may not have direct control over his present behavior; but if he had control over his drinking, then he had indirect control over his present behavior. The distinction between direct and indirect control is of great significance to ascriptions of moral responsibility, as I will explain in Sections IV and VII, where I will also address its relevance to moral obligation.

A fifth distinction is between *regulative* and *guidance* control.<sup>16</sup> One has regulative control over something just in case one can (in the relevant personal sense of 'can') bring it about and can also refrain from bringing it about. One has guidance control over something just in case one can bring it about, even if one cannot also refrain from bringing it about. This distinction has been brought to light in recent discussions of Frankfurt-type cases that seek to impugn the Principle of Alternate Possibilities (the principle that moral responsibility requires regulative control). The relevance of the distinction to both moral obligation and moral responsibility is an issue that I will address in Section V.

The five distinctions just drawn are not the only ones that can be drawn. (For example, there is a distinction to be drawn between control

---

15 Cf. Alfred R. Mele, 'Agents' Abilities,' *Noûs* 37 (2003), 448-9, on the distinction between having a 'simple' ability and being able to do something 'intentionally.' Cf. also Elinor Mason, 'Consequentialism and the "Ought Implies Can" Principle,' *American Philosophical Quarterly* 40 (2003), 321. The claim that you cannot open the safe intentionally is controversial. I base it in part on the fact that you have no justification for believing that you have a 'realistic' chance of success if you attempt to open it. Such a connection between intentional action and belief has been challenged. (See Kirk Ludwig, 'Impossible Doings,' *Philosophical Studies* 65 (1992) 257-81.) If the challenge succeeds, it remains the case that lack of 'know-how' entails lack of a certain kind of control; it is just that such control must be accounted for in some other way.

16 These terms, introduced by John Martin Fischer, have become well known. See his *The Metaphysics of Free Will* (Oxford: Blackwell 1994), 132-4 and 204. See also Zimmerman, *Essay*, 32-3, in which I used instead the terms 'standard' and 'curtailed.'

that one can exercise as an *individual* and that which one can exercise as a *member of some group*. This, too, seems morally significant.) But I will leave all other distinctions aside. It is important to note that the distinctions cut across one another, in that partial control may be either basic or comprehensive, basic control may be either deliberate or coincidental, deliberate control may be either direct or indirect, and so on. In any particular case, then, whether someone has control over something is a matter of some complexity.

### III

Let us now attend to the distinction between deliberate and coincidental control.

Consider a commonplace objection to the sort of consequentialism espoused by G. E. Moore. (It is such a natural objection that many beginning students raise it. Apparently only seasoned philosophers manage to inure themselves to it.) According to this sort of consequentialism, it is one's moral obligation to perform that act (or one of those acts) among one's alternatives that would in fact have the best consequences overall. But, so the objection goes, whether one in fact manages to do this is completely fortuitous, since it is impossible to foretell in any detail what the consequences of one's alternatives would be, and this is inconsistent with its being the case that one ought, as a matter of *moral obligation*, to bring about the best consequences.

Moore's response to this objection is well known. He says that it confounds moral obligation with moral responsibility.<sup>17</sup> Consequentialism does not imply that one is *culpable* if one fails to bring about the best consequences attainable; it only implies that one thereby does *wrong*.

I of course approve of Moore's drawing this distinction, but I do not think that it is an adequate response to the objection. For reasons that I cannot go into here,<sup>18</sup> I think we should say that whether one is morally obligated to do something depends in part on the *foreseeable*, rather than the *actual*, consequences of so acting. Moreover, for reasons that I will give in a moment, I think we should say that whether one is morally obligated to do something depends in part on what one can *intentionally* do. These two points place a strict limit on the extent to which luck can

---

<sup>17</sup> See Moore, *Ethics*, 81-2.

<sup>18</sup> The case is made, convincingly in my view, by Frank Jackson in his 'Decision-theoretic Consequentialism and the Nearest and Dearest Objection,' *Ethics* 101 (1991) 461-82.

be a factor in our fulfilling our obligations, something that seems intuitively required if we are to think of them as *moral* obligations.

Consider the safe. Suppose that there is good reason (morally speaking) for you to want to get your hands on its contents — it contains some medicine that is needed to save someone's life. Suppose that the best thing you could do under the circumstances would be to save this person's life. The consequentialist is then committed to saying that this is what you are morally obligated to do.<sup>19</sup> But this seems absurd. The absurdity is reduced, but it is surely not eliminated, by the observation that you would not be to blame for not saving the person's life. How can you be morally obligated to do something of this sort, when your managing to do so would be a matter of sheer luck? (I recognize that rhetorical questions do not have the force of arguments. Nonetheless, this question seems to me to have considerable force.) The kind of luck at issue, of course, is that which has to do with something's being beyond one's deliberate control. It has been granted that you have coincidental control over opening the safe, and hence over saving the person's life. It is deliberate control over doing so that you lack.

I am suggesting, then, that to be morally obligated to do something one must have deliberate control over doing it; if one cannot do it intentionally, then one cannot be morally obligated to do it in the first place. (This claim is actually too sweeping. I will qualify it in the next section.) If this is the case, then the *fulfilment* of obligation cannot in *this* way be subject to luck. It is important that this claim not be confused with two others. First, it does not follow that the *incurring* of obligation must be in one's deliberate control. Suppose that  $p$  implies  $q$ . Then arguably, for certain kinds of control, control over  $q$  implies control over  $p$ . Let  $p$  be the proposition that Joe is obligated to do some act  $A$  and  $q$  be the proposition that Joe can intentionally do  $A$ . Then, arguably, Joe's being in control of  $q$  implies his being in control of  $p$ . But why think that he is in control of  $q$ ? His being in control of  $A$  does not imply his being in control of  $q$ . (Recall Jean-Paul Sartre's thesis that we are 'condemned' to be free.<sup>20</sup>) Thus I see no reason to think that Joe's being in control of  $A$  implies that he is in control of his being obligated to do  $A$ .

Second, the claim (to be qualified in the next section) that one is morally obligated to perform an action only if one can intentionally

---

19 Moore himself might demur, saying that his theory concerns only what he calls 'voluntary acts' (see Moore, *Ethics*, 5-7). I do not think that this by itself will avoid the present problem, but I have no room to pursue the point here.

20 Jean-Paul Sartre, *Being and Nothingness*, trans. Hazel E. Barnes (New York: Washington Square Press 1966), 567

*perform that action* (and thereby fulfil one's obligation) does not imply the claim that one is morally obligated to perform an action only if one can intentionally *fulfil that obligation*. Whether we should accept this further claim is unclear to me. I suspect that we should not. It implies that, to have a moral obligation, one must understand the concept of moral obligation. On pain of circularity, this cannot be an analytic truth. It might be a synthetic one, I suppose, but I know of no reason to think that it is. If it is not in fact true, then one can fail to have deliberate control over fulfilling one's obligation, even if one cannot fail to have deliberate control over doing that which in fact constitutes the fulfilment of one's obligation. In this way, moral obligation is not entirely immune to luck.

Recall the second assumption noted in Section I (when restricted to culpability): one is culpable for having done something if and only if one had an obligation not to do it but did it nonetheless. If that assumption were true, then we should expect the points that I have just made about the relation between luck and the fulfilment of obligation to apply *mutatis mutandis* to the relation between luck and the incurring of responsibility. That is, we should expect, first, that one is culpable for performing an action only if one could perform it intentionally, but, second, that it is not the case that one is culpable for performing an action only if one could intentionally incur culpability for performing it. However, if the second assumption is false, as I believe, we cannot directly infer these claims about culpability from the earlier claims about obligation. Despite this fact, I think that the claims about culpability do hold nonetheless (although the first needs qualification, as does the counterpart claim about obligation). But I will not try to establish the symmetry between obligation and responsibility on this score. Instead, let me turn to a consideration of respects in which there is good reason to suspect an asymmetry between the fulfilment of obligation and the incurring of responsibility as they relate to luck.

#### IV


Consider the distinction between direct and indirect control. I said that this is of great significance to ascriptions of moral responsibility. This is because what is merely in our indirect control can at best be something for which we are merely indirectly responsible, and such responsibility can be said to be 'essentially empty.'<sup>21</sup> (One is indirectly responsible for

---

21 See Zimmerman, *Essay*, 54-61.

something just in case one is responsible for it by way of being responsible for something else.) Thomas Nagel asks, ‘How is it possible to be more or less culpable depending on whether ... a bird [gets] into the path of one’s bullet?’<sup>22</sup> The answer is that this is not possible. Consider two would-be assassins, one of whom is successful and the other of whom, due to the fortuitous intervention of a passing bird, is unsuccessful. We can depict their cases in somewhat simplified fashion as follows. Let *a* be the decision to shoot, *b* the retraction of the finger, *c* the movement of the trigger, *d* the firing of the gun, *e* the flight of the bullet, *f* the bullet’s penetration of the victim’s body, and *g* the victim’s death. Then the case of the successful assassin looks like this:

*Successful Assassin*




The arrows represent causation. *A* (which consists simply of *a*) is the minimal action (if it can be called an action at all) constituted by the assassin’s decision. *B* (which consists of *a*’s causing *b*) is the action of the assassin’s moving his finger. (In von Wright’s terminology, *b* is the ‘result’ of *B*.<sup>23</sup>) *C* (which consists of *a*’s causing *c*) is the action of the assassin’s pulling the trigger. *D* is the action of his firing the gun. *E* is the action of his causing the bullet to fly. *F* is the action of his shooting the victim. *G* is the action of his killing the victim. Note that, in light of the causal connections between *a* through *g*, the assassin may be said to do *G* *by* doing *F*, *F* *by* doing *E*, and so on. The case of the unsuccessful assassin is, by comparison, truncated. It looks like this:

---

22 Thomas Nagel, ‘Moral Luck,’ *Proceedings of the Aristotelian Society*, suppl. vol. 50 (1976), 143

23 Georg Henrik von Wright, *Explanation and Understanding* (Ithaca: Cornell University Press 1971), 66

*Unsuccessful Assassin*

(Of course, there is more to this second case, since the bullet does not simply evaporate. But I am assuming that what is left out is morally irrelevant, in that nothing further occurs for which the agent may plausibly be said to be culpable.)

Let us assume that the successful assassin is culpable for each of *a* through *g* (or *A* through *G* — it does not matter for present purposes whether we talk in terms of actions or their results). Clearly, the same is not true of the unsuccessful assassin, since he is culpable only for *a* through *e*. Thus it cannot be said that they are culpable for the same things. On the contrary, the successful assassin is *culpable for more things* than the unsuccessful assassin. Nonetheless, if two agents behave in the same way but their behavior, through luck, has different results, then there is nothing to distinguish between *them*, morally speaking. Thus the successful assassin is not *more culpable* than the unsuccessful assassin. This kind of comparison between cases reveals a key distinction between what may be called *scope* and *degree* of culpability.<sup>24</sup> The assassins bear culpability to the same degree, although the successful assassin's culpability has greater scope.

As I have depicted their cases, the assassins have direct control over *a* but merely indirect control over the ensuing events. This is because whether the ensuing events occur is not in their control at all, except insofar as *a* itself is. As it is sometimes put, whether the ensuing events occur is not up to them but 'up to nature.'<sup>25</sup> It should be clear that, just as nature can fail to cooperate by interrupting the sequence of events between *e* and *f*, as in the case of the unsuccessful assassin, so too it could interrupt this sequence at any point after the occurrence of *a*. (It could also intervene *prior* to *a*, but that raises issues that I will address in Section VII.) Interruption at any of these points would not serve to diminish the

<sup>24</sup> See Zimmerman, 'Taking Luck Seriously,' 560.

<sup>25</sup> See Donald Davidson, *Essays on Actions and Events* (Oxford: Clarendon 1980), 59.

degree of the agent's culpability, although it would of course reduce its scope. How one is to be judged, from the moral point of view, is thus entirely a function of the degree, and not of the scope, of one's culpability.<sup>26</sup> That over which one has merely indirect control is relevant only to the scope, and not the degree, of one's culpability. Thus it is irrelevant to how one is to be judged.

The emptiness of indirect culpability — indeed, of indirect moral responsibility generally — should come as no surprise. It is an example of a phenomenon that is replicated elsewhere. Consider the distinction between intrinsic and extrinsic value. One type of extrinsic value is instrumental value.<sup>27</sup> Roughly, something has instrumental value insofar as it is conducive to something else that has intrinsic value. Suppose that pleasure is intrinsically good; then, insofar as my giving you a gift is conducive to your pleasure, it is instrumentally good. But instrumental goodness, like all extrinsic goodness, is empty; it does not add to the sum total of value in the world. If it did, then I could make things better by turning my gift-giving into a Rube Goldberg-type of affair (rather than simply presenting you with the gift, I do so by doing Z, which I do by doing Y, which I do by doing X, and so on back up the alphabet to A), and that is patently absurd. Extrinsic value, being merely indirect, is essentially empty. If I give you the gift in an elaborate way, I will have done *more good things* than if I give it to you in a simple way, but I will not have done *more good*.<sup>28</sup>

---

26 At least, this is so as far as it is judgments of moral responsibility that are at issue. Agent-based evaluations that have instead to do with virtues and vices are a different matter. See Zimmerman, 'Taking Luck Seriously,' 554-5 and 569-70.

27 The term 'instrumental value' can also be used to refer to a type of intrinsic value. (See Toni Rønnow-Rasmussen, 'Instrumental Values — Strong and Weak,' *Ethical Theory and Moral Practice* 5 (2002) 23-43.) I make no such use of it in this paper.

28 There is another way in which doing *more good things* does not entail doing *more good*. Suppose that knowledge, as well as pleasure, is intrinsically good, and that, in doing some act A, John brings about both a minor instance of knowledge and a minor instance of pleasure, but that, in doing some other act B, Jane brings about a major instance of pleasure. John does (or brings about) more good things than Jane, but Jane does (or brings about) more good than John. This case is quite different from the case of gift-giving just presented. It is only extrinsic goods that are essentially empty. When it comes to intrinsic goods, the more, the better — even if one good (such as that brought about by Jane) can outweigh two goods (such as those brought about by John).

We can couch the case of John and Jane in terms of reasons. John has *more reasons* to do A than Jane has to do B, yet Jane has *more reason* to do B than John has to do A. This does *not* point to any essential emptiness in the nature of reasons. Nonetheless, I believe that there *is* such a phenomenon: we can distinguish between direct



Just as we can distinguish between direct and indirect responsibility and between direct and indirect value, we can also distinguish between direct and indirect obligation. (One is indirectly obligated regarding something just in case one is obligated regarding it by way of being obligated regarding something else.) And just as indirect responsibility and indirect value are essentially empty, so too is indirect obligation. Consider two doctors who are obligated to restore their patients' health by giving them an injection. In one case, the injection is administered simply; in the other, it is administered by means of some elaborate Rube Goldberg-type of mechanism. The first case can be depicted as follows (where *a* is the decision to inject, *b* the movement of the finger, *c* the depression of the syringe, *d* the medicine's entering the patient's body, and *e* the restoration of health):

*First Doctor*

$$a \rightarrow b \rightarrow c \rightarrow d \rightarrow e$$

(I will forego depicting how actions *A* through *E* relate to events *a* through *e*.) The second case can be depicted as follows (where the *x<sub>i</sub>* are a long series of events that mediate between *b* and *c*):

*Second Doctor*

$$a \rightarrow b \rightarrow x_1 \rightarrow x_2 \rightarrow \dots \rightarrow x_n \rightarrow c \rightarrow d \rightarrow e$$

It is of course a substantive question what constitutes a direct obligation, but let us assume for the sake of illustration that the doctors' respective obligations to bring about *e* are direct.<sup>29</sup> If we assume that each of the preceding events in the sequence is a necessary means to bringing about

---

and indirect reasons, and the latter are essentially empty. (In the gift-giving case, I have a direct reason to cause you pleasure and only an indirect reason to give you a gift or to perform any of the means *A-Z* that would enable me to accomplish this.) However, as far as I can tell, the terminology of 'more reasons' vs. 'more reason' cannot be used in any natural way to capture this distinction.

<sup>29</sup> This may be a mistake. Perhaps we should say that their direct obligation consists in bringing about some further event *f*, such as a display of respect for their patients. On the distinction between direct and indirect obligation, see W.D. Ross, *The Right and the Good* (Oxford: Clarendon Press 1930), 46.

*e*, then it follows that the doctors have obligations to bring about each of these events, too.<sup>30</sup> These obligations are, of course, indirect. In fact, let us assume that they are *merely* indirect, in that neither doctor would have an obligation to bring about any of *a* through *d* were it not for his obligation to bring about *e*. It is clear that, being merely indirect, these obligations carry no independent weight. The reason for this matches the reason given earlier. If this were not so, then the second doctor's failure to bring about any of *a* through *e* (including all of the *x<sub>i</sub>*) would constitute a more serious wrong than the first doctor's failure to do so; and this surely is not the case. The second doctor's more elaborate failure will involve *more wrongs*, but it will not involve *more wrong*. The *scope* of the second doctor's obligation is greater, but the *degree* to which he is obligated is the same.

What is interesting, however, is that, despite the similarity between responsibility and obligation regarding the distinction between scope and degree, they nonetheless differ with respect to how this distinction relates to the distinction between direct and indirect control. As the cases of the assassins indicate, the direct-indirect responsibility distinction directly tracks the direct-indirect control distinction: what is merely in our indirect control is at best something for which we are merely indirectly responsible; we can be directly responsible only for that which is in our direct control. The same cannot be said regarding the distinction between direct and indirect obligation. On the contrary, as the cases of the doctors indicate, that regarding which they have a direct obligation (restoration of health) is something over which they have merely *indirect* control; that which is in their direct control (the decision to inject) is something regarding which they have a merely *indirect* obligation. What accounts for this discrepancy?

The explanation is to be found, I think, in the fact that judgments about responsibility are agent-based whereas judgments about obligation are act-based. The degree to which a person is responsible has nothing to do with, and hence is not affected by, what is 'external' to him. The degree to which a person is obligated has much to do with what is 'external' to him. If the doctors fail to restore their patients' health, this wrongdoing has essentially to do with their patients, and the seriousness of this wrongdoing has essentially to do with how the patients are affected by it. The degree, if any, to which the doctors are culpable for failing to restore their patients' health has essentially to do, not with the patients, but with what the doctors *believe* concerning the seriousness of the

---

30 See Zimmerman, *Concept*, 69-70.

wrongdoing involved in failing to restore their patients' health. This discrepancy between responsibility and obligation implies a disparity between them as they relate to luck.

The disparity is this. As I said earlier, all control (all the control that anyone ever has) is partial control. But some instances of control are more partial (that is, less close to being complete) than others. Indirect control is always more partial than direct control, because it requires the cooperation of factors (which, by hypothesis, happens to be forthcoming) outside of one's control, factors whose cooperation is not required by direct control. If the doctors are obligated to restore their patients' health, then they can do so. (Indeed, they can intentionally do so. In my view, one can be *directly* obligated regarding (or *directly* responsible for) something only if one can (or could) intentionally do it. This provides the needed qualification to the too-sweeping claims made in the last section.) But that they can do so is due only to the cooperation of certain factors beyond their control (namely, all the events subsequent to *a* in the sequence of events from *a* through *e*). The fulfilment of an obligation is thus often a matter of luck in a way in which the incurring of responsibility is not. (Indeed, I would say that this is not just often but always the case. I cannot think of plausible examples of direct obligations that are in one's direct control.<sup>31</sup>)

## V

There are other differences regarding how luck pertains to the fulfilment of obligation as opposed to the incurring of responsibility, differences whose explanation is, I believe, to be traced to the fact that judgments about obligation are act-based whereas judgments about responsibility are agent-based. Consider the distinction between regulative and guidance control. If Frankfurt-type cases succeed in impugning the Principle of Alternate Possibilities (as I believe they do), then moral responsibility does not require regulative control; it merely requires guidance control. One way of accounting for this is as follows: if an agent performs some act freely, then it does not matter, insofar as *he* is to be judged from the moral point of view, whether he could *in fact* have acted otherwise; all that matters (as far as alternate possibilities go) is that he *believed* that he

---

31 Contrast H.A. Prichard, *Moral Obligation* (Oxford: Clarendon 1949), 31-3, and W.D. Ross, *Foundations of Ethics* (Oxford: Clarendon 1939), 153-4. I think both authors are misled in part by their failure to distinguish the distinction between partial and complete control from the distinction between direct and indirect control. (They also make what strikes me as an unwarranted appeal to the knowability of obligation.)

could have acted otherwise. But moral obligation is different; it requires regulative control. The reason is this. Just as 'ought' (in the sense of 'obligated') implies 'can,' so too does 'wrong'; and since one is obligated to do something just in case it is wrong of one not to do it (that is, it is wrong to refrain from doing it), it follows that one is obligated to do something only if one can do it and can also refrain from doing it.<sup>32</sup> As far as I can tell, there is no plausible corresponding argument for the claim that moral responsibility requires regulative control,<sup>33</sup> since in this context there is no plausible counterpart to the principle just mentioned that links obligation and wrongdoing.

## VI

The principle that one is obligated to do something just in case it is wrong of one not to do it also plays a role in the question whether moral dilemmas are possible. Such dilemmas may be understood in one of two ways: either as situations in which an agent cannot avoid wrongdoing (I will call these *obligation-dilemmas*) or as situations in which an agent cannot avoid culpability (I will call these *responsibility-dilemmas*). (Writers often fail to distinguish between these two kinds of dilemma, due to a failure to distinguish between wrongdoing and culpability generally.<sup>34</sup>) In my view, obligation-dilemmas are impossible. The arguments for this claim are familiar.<sup>35</sup> Note that, if an obligation-dilemma arose over which one had no control, one would be a victim of bad luck. The luck would be bad because of the wrongdoing involved. This variety of bad luck is so bad that some have called it 'tragic.' Thus the arguments against the possibility of obligation-dilemmas are arguments against the

---

32 See Michael J. Zimmerman, 'Obligation, Responsibility, and Alternate Possibilities,' *Analysis* 53 (1993) 51-3, and Zimmerman, *Concept*, 85-9. Cf. Ishtiyaque Haji, 'Alternative Possibilities, Moral Obligation, and Moral Responsibility,' *Philosophical Papers* 22 (1993) 41-50; *Appraisability*, 53; and *Deontic Morality*, ch. 3.

33 I consider and reject other kinds of arguments for this claim in Michael J. Zimmerman, 'The Moral Significance of Alternate Possibilities,' in *Freedom, Responsibility, and Agency*, eds. David Widerker and Michael McKenna (Aldershot: Ashgate Press 2003) 301-25.

34 Cf., for example, Daniel Statman, 'Introduction,' in *Moral Luck*, ed. Daniel Statman (Albany: SUNY Press 1993), 20-1.

35 Recall from the opening paragraph of Section I that it is only overall obligation that is at issue here. Thus I am not dismissing the possibility of conflicts of prima facie obligation. Cf. Zimmerman, *Concept*, ch. 7. (Note that the arguments apply only to what I there call 'basic' dilemmas.)

possibility that life is tragic in this sense. Note also that the question whether obligation-dilemmas are possible concerns luck regarding the incurring of obligations, not their fulfilment. The proponent of such dilemmas typically concedes that, for each obligation that an agent allegedly has, the agent has control over fulfilling that obligation. The opponent of such dilemmas also typically concedes this point. The question that such dilemmas concerns is whether it is possible to incur an obligation to perform some act *A2* — an act that one can perform — if one already has an obligation to perform some other act *A1*, an act which is such that one can perform it, but only if one does not perform *A2*. The opponent of obligation-dilemmas claims that such a possibility cannot arise under any circumstances, and thus *a fortiori* cannot arise outside the control of the agent.

One argument against the possibility of obligation-dilemmas conjoins the principle that one is obligated to do something just in case it is wrong of one not to do it with the further principle that, if one is obligated to do something, then it is not wrong of one to do it. These principles imply (jointly, but not singly) that it cannot happen that one is both obligated to do something and obligated not to do it. I believe this argument to be sound. Note, however, that such an argument is not available to the opponent of responsibility-dilemmas, since, as noted in the last section, in the context of responsibility there is no plausible counterpart to the first principle.

Another argument against the possibility of obligation-dilemmas rests on the principle that 'ought' implies 'can' conjoined with the Principle of Agglomeration, according to which, if one ought to do each of two things, then one ought to do them both. I believe this argument also to be sound. But again, such an argument is unavailable to the opponent of responsibility-dilemmas, since in the context of responsibility there is no plausible counterpart to the Principle of Agglomeration.

The unavailability of such arguments to the opponent of responsibility-dilemmas does not of course mean that such dilemmas are possible. One reason to think that they nonetheless are possible is this: even if it is impossible for it to be the case that all one's alternatives are *in fact* wrong, it is not impossible for it to be the case that all one's alternatives are such that one *believes* them to be wrong; and if responsibility turns, as I have claimed, on the agent's beliefs about wrongdoing rather than on actual wrongdoing, then we cannot infer the impossibility of responsibility-dilemmas from that of obligation-dilemmas. Despite this, it might still be true for some countervailing reason that one can never be faced with a situation in which one will be culpable for whatever it is that one does. Unfortunately, however, even if this were so, it would *not* mean that responsibility-dilemmas are impossible, as I will now briefly explain.

## VII

The disparity noted in Section IV between responsibility and obligation insofar as they concern direct and indirect control can be heightened. Return to the case of the Unsuccessful Assassin. I said that the fortuitous intervention of nature in the form of a passing bird, while reducing the scope of the assassin's culpability, would not diminish its degree. But I also noted that nature could intervene earlier in the sequence of events from *a* to *e*; indeed, it could intervene even *prior* to *a*, the assassin's decision to shoot. For example, it could happen that, just as he is about to make this decision, the assassin is seized by a sudden sneeze that prevents him from making it. If the fortuitous intervention of the bird does not diminish his culpability, I cannot see how the fortuitous intervention of the sneeze could do so. In this case, though, there would be nothing *for* which he would be culpable, whether directly or indirectly. He would be culpable to some degree, but his culpability would have zero scope. He would be, as I have put it elsewhere, culpable '*tout court*.'<sup>36</sup>

The possibility of responsibility *tout court* shows that it is not only indirect responsibility that is essentially empty; direct responsibility is too! It shows, moreover, that responsibility is not essentially retrospective; if one is responsible but not responsible *for* anything, then one's responsibility lacks *any* 'perspective,' whether 'retro' or otherwise. I find no parallel to the phenomenon of responsibility *tout court* in the area of obligation, however; I can make no sense of the idea of someone's being obligated without there being anything regarding which he is obligated. Responsibility *tout court* is possible because how a person is to be judged from the moral point of view<sup>37</sup> turns, at bottom, not on how he actually behaves but on how he counterfactually would behave. If, but for some stroke of luck, the unsuccessful assassin would have behaved just as the successful assassin did, then the former is just as culpable as the latter. Obligation *tout court* is not possible, because whether an act is obligatory turns not on what options an agent would counterfactually face but on the options he actually faces. If an agent has no actual options, then he has no actual obligations; if he has an actual obligation, there must be something regarding which he is obligated.

The idea of responsibility *tout court* is in a way paradoxical. One arrives at it by taking seriously the idea that control is essential to

---

36 See Zimmerman, 'Taking Luck Seriously,' section IV, for a fuller rendition of this argument, and sections V and VI for an elaboration and defense of it.

37 For a necessary qualification, see note 26 above.

responsibility: we cannot be responsible for what is not in our control. Yet the upshot is that we can be responsible even if we lack control over what we do; all that matters is (roughly) whether we *would* act in a certain way if we *did* have such control. In this way, we can be 'condemned' to be responsible. (This is once again reminiscent of what Sartre says about being 'condemned' to be free. In pointing this out, though, I do not mean to be taking on any excess existentialist baggage.)

It might seem that, just as one can be condemned to be responsible, so too one can be condemned to be obligated. But there is a crucial difference. It is true that, once one has reached a certain level of maturity, it can happen that one cannot control whether one *has* certain obligations. (You might think that this is false, insofar as someone might avoid having any obligations at all — by committing suicide, for instance. But not so fast! One may have an obligation not to avoid having obligations.) Thus, whether one *is* obligated can be a matter of luck. But this has only to do with the *incurring* of obligation. It does not follow that whether one *fulfils* the obligations that one has can be a matter of luck. On the contrary, given that obligation requires regulative control (as argued in Section V), whether one fulfils one's obligations is (in this way, and to this extent) *not* a matter of luck. Thus, even if one can be condemned to *be obligated*, one cannot be condemned to *do wrong*. By contrast, the possibility of responsibility *tout court* implies that one can indeed be condemned to *be culpable*. I am afraid, therefore, that there is an important sense in which life can after all be tragic. Indeed, I would call this Tragedy with a capital 'T.'

But not to worry! Life can also be Blessed, with a capital 'B.' That is because the possibility of responsibility *tout court* implies that one can be 'condemned' to be laudable. In fact, life can be downright Bewildering, also with a capital 'B.' That is because it can be both Tragic and Blessed at once.

## VIII

Let me summarize the findings of this paper.

I drew five distinctions concerning control. I put the first distinction (that between partial and complete control) aside. This was because nothing is ever in our complete control, and I assumed that neither obligation nor responsibility requires such control.

I also put the second distinction (that between basic and comprehensive control) aside, not because this distinction is not relevant to questions about obligation and responsibility, but because luck is a function only of basic, and not also of comprehensive, control.

The third distinction (that between deliberate and coincidental control) is important to both obligation and responsibility. In my view, both obligation and responsibility require deliberate control. More precisely, one cannot be either directly obligated regarding or directly responsible for something without its being the case that one has deliberate control over that thing. In this way, both obligation and responsibility preclude luck. However, it remains possible that one not have deliberate control either over fulfilling one's obligation or over incurring responsibility, and so luck can play a role in both areas in this way. As far as I can tell, no disparity between obligation and responsibility arises in connection with this third distinction.

But the fourth distinction (that between direct and indirect control) is different. Here a disparity does arise. Insofar as direct responsibility requires direct control but direct obligation does not, luck can be a factor in the fulfilment of obligation in a way in which it cannot be in the incurring of responsibility.

The fifth distinction (that between regulative and guidance control) also reveals a disparity between obligation and responsibility, inasmuch as obligation requires regulative control but responsibility apparently does not.

There is, moreover, a discrepancy between obligation and responsibility with respect to the sort of bad luck involved in facing a moral dilemma, inasmuch as obligation-dilemmas are impossible but, unfortunately, responsibility-dilemmas are not.

The net result of my survey, then, is as follows. With respect to the incurring of obligation: this can (but, of course, need not) happen entirely beyond the agent's control; however, an agent cannot be faced with obligation-dilemmas. With respect to the resolution of obligation: this cannot be beyond the agent's partial, regulative control; and insofar as it is, in particular, direct obligation that is concerned, its resolution cannot be beyond the agent's deliberate control. Finally, with respect to the incurring of responsibility: inasmuch as it is responsibility *for* something that is concerned, the incurring of responsibility cannot be beyond the agent's partial, guidance control; and insofar as it is, in particular, direct responsibility for something that is concerned, its incurring cannot be beyond the agent's deliberate, direct control. However, responsibility *tout court* may be incurred in a way that is wholly beyond the control of the agent.

A more detailed investigation of the many and various connected issues that I have broached in this paper might of course result in different findings. However, even at this point it seems safe to conclude that, although there are several ways in which the relation between luck and obligation is similar to the relation between luck and responsibility,



there are also several ways in which the relations are dissimilar. We would do well, therefore, not to confuse the two.<sup>38</sup>

---

38 Thanks to members of audiences at the Universities of Reading, Florida, and Calgary for helpful comments. Thanks in particular to John Baker, Randy Clarke, David Copp, John Cottingham, Ish Haji, Bob Kane, Noa Latham, Kirk Ludwig, Al Mele, Mark Migotti, Dennis McKerlie, David Oderberg, Sigrun Svavarsdottir, Jon Tresan, and Andrew Williams. Thanks also to the referees for *Canadian Journal of Philosophy*. Thanks finally to the National Endowment for the Humanities for support while this paper was written.