



Moral Responsibility and "Moral Luck"

Author(s): Brian Rosebury

Source: *The Philosophical Review*, Oct., 1995, Vol. 104, No. 4 (Oct., 1995), pp. 499-524

Published by: Duke University Press on behalf of Philosophical Review

Stable URL: <https://www.jstor.org/stable/2185815>

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.

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



Duke University Press and are collaborating with JSTOR to digitize, preserve and extend access to *The Philosophical Review*

JSTOR

Moral Responsibility and “Moral Luck”

Brian Rosebury

This paper has two purposes. The first (part 1) is to defend a distinctive account of moral responsibility; the second (part 2), to argue that “moral luck,” understood as a susceptibility of moral desert to lucky or unlucky outcomes, does not exist. The strategy will be to show that if moral responsibility is correctly understood, the phenomena that supposedly require “moral luck” as part of their explanation are fully explicable without it. The argument will turn upon the proposition that epistemic inquiry is an indissoluble part of moral responsibility in a world of intellectually and morally fallible persons, and that judgment on the moral decision making of others not only must, but in principle can, adjust for this fact.

1. Moral Responsibility

Analyses of moral decision making generally proceed with the aid of test cases in which the relevant circumstances are precisely defined and, by implication, known for certain. For the purposes of these cases one is supposed to know, for example, that the runaway trolley will kill five people unless one pulls the lever to let it kill one person instead, but to be significantly unsure about whether it is morally justifiable to pull the lever. The purpose of this type of *mise en scène* is evidently to isolate the strictly moral dilemma from infection by irrelevant practical or epistemic considerations of a kind that the unphilosophical are only too liable to intrude into discussion.

Real cases of moral decision are usually quite different. When we make practical moral decisions—for example, in deciding how to bring up our children—we are generally as much exercised by questions of knowledge as by questions of value; in deciding whether to restrict children’s access to television, for example, we ask ourselves whether watching television is likely to advance or retard their educational development, influence their social behavior for better or worse, frighten them, relax them, deprive them of exercise, and so on. In morally significant cases we feel not only a desire but an obligation to get the answer right, if we can. However, we

often get it wrong, and we notice that others do the same. This epistemic fallibility is itself one of the empirical facts that moral agents need to take into account in making decisions; and others who judge their decisions must take it into account also, if they are to judge fairly, distinguishing between decision making vitiated by avoidably and by unavoidably insufficient knowledge, by culpably and by non-culpably ineffective deployment of knowledge. Conspicuous failures to pay attention to the relevant circumstances of our actions, or to predict realistically their consequences, tend in fact, however virtuous our intentions, to attract a distinctively moral kind of censure. The distinction holds good between adopting a value and pursuing inquiry into the best way of acting in conformity with it, but the latter as well as the former belongs to the sphere of moral responsibility, notwithstanding the fact that such inquiry also occurs in nonmoral contexts.

If we take the welfare of others as our value, we encounter special difficulties in judging with confidence which actions will actually serve their welfare, and whether we ourselves are capable of those actions. Freud and Proust have taught us, if we needed any teaching, not only the partial, perhaps radical, unknowableness of other people, but the contribution our psychological attitudes towards them make to that unknowableness. We see in others what we wish (or fear) to see; we attribute to them the emotions that our own emotional life disposes us to find natural and intelligible. We cannot necessarily sidestep this danger just by complying with the other person's expressed preference, supposing that we learn of it, for we remember that we ourselves are often mistaken about our own interests, and sometimes, whether deliberately or by inadvertence or through error, misrepresent them to others. Moreover, when the acts that express our good intentions are observed by someone else, the influence of self-deception or avoidable ignorance or misjudgment is often very apparent (unless it is the observer, no less fallible than the observed, who is mistaken). The relatives of an elderly, occasionally vague invalid tend him with a well-intentioned supervision of his every move that an outsider can recognize as oppressive and resented, while an exaggerated estimate of his mental deterioration is taken to justify disregarding his wishes, complaints, conversational gambits. An affectionate but lethargic parent oscillates between extremes of leniency and strict-

ness. A civilized politician in the grip of wishful thinking appeases an insatiable aggressor.

These familiar examples reinforce the point that moral obligation is not satisfied simply by commitment to a value, such as the welfare of others. If we are to begin trying to make such a commitment effective we must engage in an often difficult assessment of empirical circumstances and probabilities, an assessment that will have to take account of our own fallibility of judgment, and of the limits of our practical capabilities. This reflexive consideration is indeed widely acknowledged in philosophy as well as in commonsense moral decision making. Part of the case for rule- as against act-utilitarianism rests on a recognition of the limitations of the individual personality and the solitary perspective confronted with each new occasion for judgment: the rule takes precedence over rule-free judgment in individual cases because, as Berkeley puts it, "the best men, for want of judgement, and the wisest, for want of knowing all the hidden circumstances and consequences of an action, may very often be at a loss how to behave themselves, which they would not be in case they judged of each action by comparing it with some particular precept, rather than by examining the good or evil which in that single instance it tends to procure; it being far more easy to judge with certainty, whether such or such an action be a transgression of this or that precept, than whether it will be attended with more good or ill consequences."¹ The rule is designed to minimize the effect of individual fallibility and restricted knowledge.

Let us assume, for the purposes of the rest of this paper, the most broadly stated altruistic morality: the moral agent, we suppose, tries to act in such a way as to serve the interests of other persons and judges others against the same criterion. He may employ a utilitarian calculus or he may not, but so far as the choice of value is concerned he has chosen well by choosing the well-being of others. We will also simplify the discussion of examples by discounting the agent's own interests, except insofar as they subserve indirectly those of others: from a moral point of view, I aim (in the normal case) to preserve my life, health, soundness of judgment and emotional stability in order to go on doing useful work, caring for my children, interacting supportively with others, dis-

¹Berkeley, *Passive Obedience*, 3d ed. (1713), 9.

charging civic responsibilities, and so on. (That there is scope for self-interested self-deception here, for awarding myself special comforts and recreations today in order to make myself a *supremely* fit and efficient altruist tomorrow, is precisely one of the kinds of insight the paper will seek to emphasize.)

As an altruist, then, I interrogate my knowledge, including my self-knowledge, in order to make my altruism as effective as possible. That there is an obligation to consult my available relevant knowledge in making even the most nearly instantaneous of morally significant decisions seems intuitively clear. Consider, for example, near-instantaneous decisions to vary the speed of an automobile in motion. I know such facts, or probabilities, as that if I drive at eighty miles per hour through a city after dark I increase the likelihood of killing or injuring someone. If I drive without taking account of such knowledge I am morally culpable, notwithstanding my good intention not to hit anyone. Some epistemic context is entailed by the very concept of a decision: if no relevant state of affairs is predicated by me there is no decision to act, just a reflex, and no question of moral responsibility can arise for a reflex *per se* (though there may be responsibility for anterior decisions that cause me to have a reflex when I might otherwise have been able to make a decision, such as an anterior decision to drive so fast as to deprive me of the possibility of making certain decisions later.)

Further initiatives of inquiry and self-interrogation may be appropriate, in light of the fact that few people have comprehensive knowledge about a situation of any complexity. The basic assumption is that for any contemplated action (or deliberate abstention from action) there is a minimum consultation of relevant knowledge below which the action (or inaction) is *prima facie* irresponsible. It must be at least possible in many cases that this minimum is greater than the knowledge I have immediately to hand: hence there may be an obligation to acquire additional knowledge, or, where an item of relevant knowledge is unattainable, to weigh carefully the significance of this shortfall for my project. In addition, there is a plausible presumption of universal human moral and intellectual fallibility that bears upon the self-interrogation itself. (People whose behavior or speech suggests that they regard themselves as exceptions on either score are prone to a disapproval from others that has a clearly moral tinge.) I need to consider especially

the possible impact of special intellectual factors (e.g., a poor track record in judging character) or special moral factors (e.g., a vice, passion, or vested interest) in distorting my judgment. It makes moral sense to bring the inquiry to an end when the likely increment of altruistic value achieved by pressing the questioning further begins to be outweighed by the likely increment achieved by acting now, or by turning to another project that perhaps merits longer reflection. (Indeed there would be circumstances in which too much questioning would become culpable.) In practice, most people in most circumstances will conclude the self-catechism fairly quickly, and no doubt they are right to do so given the dangers of a disabling perfectionism.

The assessment of another's actions is itself a morally significant action and therefore the interrogative sequence again comes into play. The knowledge available to and deployed by the initial agent becomes itself an object of inquiry for the author of the judgment. We know that even an agent endeavoring to act altruistically and to gather and deploy relevant knowledge with appropriate vigor will from time to time act with counter-altruistic results because of the epistemic shortfall. To pass fair moral judgment on an action we must, therefore, take into account the knowledge, as well as the value, that justifies it to the agent; and since knowledge varies from moment to moment we must direct our judgment of the decision to a particular moment. We should not blame an agent who, granted the evidence available, did as he ought, though subsequent evidence would have led him to a different view of what he ought to do.

This is not to say that in such cases I myself am entitled to accept, without further self-questioning, such exculpation as might be granted by a reasonable observer. On the contrary, I should view with more than average suspicion the epistemic grounds for self-exculpation that occur to me or are proposed to me by others, since I have a potentially distorting interest in believing them. Moreover, my imperfect access to knowledge, and my participation in the universal human fallibility in its pursuit and deployment, are not considerations to be cited *by me* in mitigation of my imperfect success in making altruism effective; rather, these deficiencies represent dangers I need to allow for in the process of decision making. For there is a difference between (typically private) self-judgment and (typically public) judgment of others. When A privately

judges A her moral objective is to improve, by a discipline of self-criticism, her efforts to serve the interests of others: no question of comparison with others arises, and in a sense perfection is the target. When A publicly judges B, comparison is of the essence: to serve their ethical purpose adequately such public judgments must respect a collectively viable norm that ensures that the language of moral praise or blame does not become devalued, by either a stringency or a laxness in its use that comes to negate its function of discriminating among acts across the whole relevant population. (The situation is complicated by the fact that alongside self-judgment we can try to imagine how another might judge us, and in addition to judging others against the collective norm we can try to imagine how we would judge ourselves if we were they.)

Correspondingly, and with the reservation just made, our judgment of others' perseverance with epistemic inquiry in support of moral decision making must be consistent with a collectively viable norm: to avoid justified blame from others, an agent must take what we typically call "reasonable" steps to foresee the consequences of his action for others. What is reasonable will be understood more stringently the more serious the issues at stake. The force of these requirements will, to maintain fairness in judging the acts of different persons, be adjusted to take account of defective powers of reasoning, special obligations to deploy knowledge attaching to particular roles and offices, and so forth. But the requirement of pursued and deployed knowledge, however it may vary in severity, is part of the moral obligation itself: there is a point at which an agent can no longer defend himself against moral blame with the reply that he did not foresee the outcome of his action or was unaware of certain relevant circumstances. This explains why, for example, accomplices and defenders of Hitler or Stalin who believed them, sincerely but mistakenly, to be benefactors of humanity may be liable to moral reprehension, not merely to a low mark for intellectual performance. If it is not that they were simply less intelligent than the others who acted differently, if it can be shown that they failed to pursue or deploy relevant knowledge with the energy that (given their capacities and circumstances) would have been appropriate in view of the gravity of their actions or the actions with which they associated themselves, a gravity that increases the stringency of the requirement of deployed knowledge, then they are morally culpable. It also explains why

certain dilemmas instanced in moral philosophy, such as that in which the agent knows that any number of prisoners will be spared from death by the terrorist or fascist police chief if only she herself will shoot one prisoner, strike us as falsely posed: the empirical consideration that the agent could not know for sure that the promise would be honored is not a complication that can be discounted for the sake of the ethical clarity of the example, but an irremovable part of the moral predicament the agent must resolve.

2. Moral Luck

In section 1 I argued that moral responsibility incorporates an obligation to pursue and deploy knowledge as well as to adopt a value, and that moral judgments upon action must take account of the knowledge available to the agent, and the agent's deployment of that knowledge, at the point of decision. It will now be convenient to develop this view through a critique of the notion of "moral luck," introduced in an exchange of papers between Bernard Williams and Thomas Nagel in 1976 and extensively debated since.² I maintain that "moral luck" proper, that is luck in the *deserving* of praise or blame (as distinct from luck in the *receipt* of expressed praise or blame), is a chimera, and that a recognition of the role of epistemic fallibility as a variable in moral decision making will help us to dispose of it.

Williams and Nagel concur in claiming to show that the intuitively plausible principle that moral value is immune to luck, that we can justly be morally blamed or commended only with respect to performances or outcomes within our control, is unsustainable on

²*Proceedings of the Aristotelian Society*, supp. vol. 50; papers reprinted in revised form in B. Williams, *Moral Luck* (Cambridge: Cambridge University Press, 1981), 20-39, and T. Nagel, *Mortal Questions* (Cambridge: Cambridge University Press, 1979), 24-38. Various criticisms and modifications of their arguments are offered in H. Jensen, "Morality and Luck," *Philosophy* 59 (1984): 323-30; J. André, "Nagel, Williams and Moral Luck," *Analysis* 43 (1983): 202-7; H. Smith, "Culpable Ignorance," *Philosophical Review* 92 (1983): 543-71; N. Richards, "Luck and Desert," *Mind* 95 (1986): 198-209; S. Sverdlik, "Crime and Moral Luck," *American Philosophical Quarterly* 25 (1988): 79-95; M. J. Zimmerman, "Luck and Moral Responsibility," *Ethics* 97 (1987) 374-86; D. Statman, "Moral and Epistemic Luck," *Ratio* 4 (1991): 146-56; B. Browne, "A Solution to the Problem of Moral Luck," *Philosophical Quarterly* 42 (1992): 345-56.

reflection. For Nagel this leaves us with an incoherence in our moral conceptions; for Williams it also exposes the unsustainability of morality's claim to cohere with rationality. There is more to Williams's further argument than can be explored here, including much with which I agree but which is not, in my view, dependent on the notion of "moral luck." It is the first and basic claim, and its threat to a coherent view of moral responsibility, that I mainly wish to challenge.

There are two types of apparent difficulty for a luck-free moral scheme. At the level of intuitive moral judgments, a variety of examples seem to suggest that our moral evaluation of another's conduct is fairly commonly influenced by factors beyond his or her definite knowledge at the time of decision. Williams gives the example of the painter Gauguin, abandoning his family in order to develop his art untrammelled by other responsibilities. It can be argued that this action is justifiable retrospectively if the agent succeeds in becoming a great painter, and unjustifiable if he fails: but this outcome cannot be known at the time of decision and is largely dependent on contingencies that postdate decision. In such a case, according to Williams, one finds oneself passing judgment, not on an instantaneous decision in the light of the evidence then available to the agent, but on the success or failure of a morally significant project pursued over a period of time and susceptible to various kinds of luck. In another example, Nagel observes that if one negligently leaves the baby in the bath with the tap running, "one will realize, as one bounds up the stairs towards the bathroom, that if the baby has drowned one has done something awful, whereas if it has not one has merely been careless."³ In this case, the decision (to leave the baby unattended for a while) is the same in both cases, founded on the same poor accomplishment in the deployment of knowledge, yet the moral verdict differs greatly according to an outcome determined by luck. Again, the law punishes murder more severely than attempted murder, even where the criminal has nothing better than his own incompetence to thank for the survival of his victim. In all these cases one has the option of repudiating the intuitive or juridical view in favor of the Kantian doctrine that moral praise or blame can only properly attach to the agent's willing and ought to be indifferent to performed ac-

³Nagel, *Mortal Questions*, 30-31.

tions and their outcomes: but both Nagel and Williams suggest that the examples show this view to be intuitively unacceptable.

There is also a more general problem expressible through the notion of "moral luck." For a luck-free moral scheme we need a workable distinction between "agency," which is the object of judgment, and the operation of factors beyond an agent's control. But the former seems to be progressively impoverished as, on reflection, we withdraw from it, not only lucky and unlucky consequences of action, but whatever conditions of action are determined by the agent's immediate circumstances and historical opportunities; finally it is reduced to a vanishingly small point by the consideration that, as Williams puts it, "the dispositions of morality, however far back they are placed in their direction of motive and intention, are as 'conditioned' as anything else."⁴ Here again one is apparently driven to a choice between, on the one hand, the admission of luck into the grounds for moral judgments, and on the other, allowing something as metaphysically problematic as Kant's unconditioned will to replace the action as the object of moral judgment.

I hope to show that the notion of "moral luck" can be deprived of most of its plausibility by a close examination of the examples of intuitive and legal judgment that supposedly substantiate it; the examination will simultaneously serve to display the explanatory utility of the idea of moral responsibility presented in section 1. But before turning to this analysis, I will present and defend some fundamental claims that were implicit in section 1 but now need to be spelt out in order to make sense of the more general problem about judgment and responsibility.

The appropriate primary "target" of moral judgment I take to be the performance of an intentional act in the light of a certain amount of knowledge. (In the category of "acts" I am including for this purpose all aspects of intentional behavior, including "manner," "attitude," and the like: I am not countenancing a definition of 'act' that would embrace, for example, the mechanical performance of paying a debt or offering a lift to the hospital, but exclude the grudging or compassionate manner of the performance.) If this is correct, it has two consequences relevant to the debate about "moral luck." The first is that judgment is always of

⁴Williams, *Moral Luck*, 21.

an act, at a given time, given the agent's knowledge and reasonable beliefs at that time. Part of the significance of this point is that it lends support to the intuition that there is no contradiction in welcoming happy outcomes of wrong actions, or deploring unhappy outcomes of right actions. We can believe that it was a good thing that there was a Democratic administration in 1960–64, believe that there would not have been a Democratic administration had there not been electoral irregularities, and yet still believe without inconsistency that the authors of the irregularities were wrong to commit them. Our judgment of an act does not constrain us to an evaluation of the actual consequences of that act; nor does our evaluation of the actual consequences constrain our judgment of the act. Courts of law recognize this principle in considering what a person might reasonably have expected to be the consequence of her act, irrespective of the actual outcome. Its relevance to Williams's Gauguin example will soon become clear.

The second implication is that feelings or dispositions are the objects of secondary (derivative or generalized) judgments only, as the sources or consequences of acts: that they lie outside the scope of moral judgment if they are disconnected from acts. The relevance here is that, assuming that we do not choose our dispositions, the case for "moral luck" seems to be clinched if (unchosen) dispositions are independently the objects of moral judgment.⁵ (Williams's more radical point, of course, is that we do not choose our acts either, just because we do not choose what causes them. But this consideration either cancels through, so to speak, or dissolves all moral judgment of persons, including any judgments that might substantiate "moral luck.")⁶ It is no surprise, then, that Na-

⁵But compare the arguments of Richards and André, who both seek to *defend* common moral notions against the subversive threat of lucky factors that bear upon acts by subordinating acts to dispositions as objects of moral judgment. "In its central sense . . .," André puts it, "morality refers to excellence of character" (205); and qualities of character can be judged (morally) without its being contradictory to add that a person is not to blame for having a bad character. Our moral notions are partly—perhaps centrally—Aristotelian (concerned with whether a person's character is to be admired and emulated) and only partly Kantian (concerned with right intention). However, this view seems to complicate rather than resolve the problems Nagel and Williams identify, creating disjunctions between character and blameworthiness to add to those between blameworthiness and outcome.

⁶A rare case of a philosopher who maintains, in effect, that we do choose our dispositions but do not choose our acts, and draws the appropriate

gel maintains, against Kant, that a *disposition* can appropriately be blamed—or, presumably, praised—even where the agent's will conclusively overrides it so far as actions are concerned. An envious person who controls his impulses and behaves entirely non-enviously "still has the vice" of envy, as Nagel puts it, and is morally condemned for it.⁷

Nagel offers this last statement as a blunt fact about our blaming practices. But an appeal to intuition or to practice seems to have less than its full value if, as here, one of the empirical conditions of actual moral judgments—the imperfect knowledge enjoyed by those who assess as well as by those who act—is suppressed. What in normal conditions of judgment we condemn under the name of a "vice" is not an unacted-on temperamental inclination (how could we know of this in the normal case?), but a pattern of actions that makes itself apparent and that it is natural to construe as the external correlative of a disposition that generates, and is perhaps reinforced by, the individual acts. In the case Nagel cites, the only person likely to know that the agent feels envy, unless he displays it through overt acts, is the agent himself; and his envious inclination presents itself to him as an obstacle he succeeds in overriding, not a cancellation of his success. (He must feel that it would be easier for him if the obstacle were absent, but this thought is not a *condemnation*.) As for others, since they cannot *ex hypothesi* know of his envy, there is no opportunity in the normal case to test Nagel's insistence that intuition—even if irrationally—condemns it. People are assessed for what they are *like*, Nagel observes: but it might more exactly be said that they are assessed for what they *show themselves* to be like by their actions (including their manner, attitudes, etc.). There is the special case in which an inclination, unacted on, comes to be known because it is confessed; and it is instructive to consider under what circumstances an agent might feel impelled to make such a confession.

In actual experience, the state of affairs postulated by Nagel, in which internal inclination and external behavior are wholly contrasted, is a rare and unstable one, and the awareness that it is unstable, that with the best will in the world (to exploit a cliché)

moral conclusions, is Schopenhauer (*Essay on the Freedom of the Will*, trans. K. Kolenda (New York: Liberal Arts Press, 1960)).

⁷Nagel, *Mortal Questions*, 33.

few vicious inclinations fail to leave at least faint traces in action, is built into our praising and blaming practices. Suppose A confesses to an unacted-on inclination to, say, cruelty, and it seems to us that she sincerely intends that it should remain unacted on. Despite our conviction of her sincere intention, we may agree with Nagel that we are likely to feel a disagreeable emotion—disapproval or condemnation—of some sort. There would appear to be (at least) eight alternative condemnatory attitudes available in a context like this.

1. I condemn all acts of cruelty (irrespective of persons).
2. I condemn the inclination to cruelty (irrespective of persons).
3. I condemn A's inclination to cruelty (though unacted on).
4. I condemn A for having an inclination to cruelty (though unacted on).
5. I condemn A's acts of cruelty.
6. I condemn A for carrying out acts of cruelty.
7. I condemn A's inclination to cruelty for causing A to perform acts of cruelty.
8. I condemn A for allowing her inclination to cruelty to lead her to perform acts of cruelty.

By 1 and 2 one reminds oneself that cruelty is a bad thing. Just thinking about cruelty should lead one to do this, so these attitudes will be evoked by A's confession; no condemnation of A is involved, however, so no question of A's "moral luck" arises. It is essential to Nagel's position that there be more to our condemnatory attitude than 1 and 2. Condemnations 5, 6, 7, and 8 are *ex hypothesi* inapplicable to A's case, and if we regard any one of these as the foundational type of moral judgment, then our belief will indeed, on the face of it, be uncomfortably at variance with any intuitive condemnatory attitude we may have towards A (though a dedicated Kantian, noting the victory of will over inclination that rescues A from condemnation 8, might genuinely admire her, in which case no conflict between doctrine and intuitions would arise.) Nagel's argument requires that in addition to 1 and 2 A's confession promote 3 and 4. At a glance, 3 looks to be entailed by 2, but the true position is more complex. The reason we unhesitatingly condemn what we call "an inclination to cruelty" is that it typically leads to cruel acts. Indeed, to assert that someone has an "inclination

nation to cruelty" is, in ordinary discourse, to imply that she *has* performed some cruel acts, and that we are inferring her temperament from those acts; we are not usually understood to be grounding such an assertion on special access to her consciousness. If, improbably, psychological research discovered that people possessing a disposition plausibly described as "cruel" (such as a disposition to entertain fantasies of themselves hurting others) in fact typically performed *fewer* acts of cruelty than those lacking that disposition, we would at any rate be puzzled over whether the disposition should be condemned, and might even be tempted to redefine it as an oblique kind of compassion. (For the psychological sources of a compassionate temperament may be complex. Perhaps acute sensitivity to the possibility of others' pain is inseparable from lively sadistic impulses. In that case moral judgment of the person, as having the "inclination to cruelty" or the "inclination to kindness," would turn on what she was motivated to *do*. We would still in some sense wish to deplore the sadistic impulses, even when they do not motivate cruel actions, even when they function indirectly to motivate kind actions, but what is left to deplore about them is what they express, that is to say, the fact that they envisage actions that, were the agent to carry them out, would be bad actions. It would not be consistent to blame the agent merely for having them if a person could not be compassionate who did not have them.) Moreover, actions are the best index we have, though a fallible one, of the comparative intensity or seriousness of a person's desires. If the desires of a particular person never led to cruel acts, they would not credibly exemplify the inclination to cruelty, but would be more plausibly described as a string of velleities within a compassionate emotional economy. (It would be even less clear how the "vice" of cruelty could be something that never, or even rarely, led to cruel actions. 'Vice' implies, as the normal case, a settled habit of vicious behavior.)

The provision in 3 that A performs and will perform no cruel acts means therefore that the crucial justification for 2 is exceptionally negated in 3, as if a particular mushroom of a species of deadly poisonous mushroom happened to be perfectly edible because it freakishly secreted an antidote. In fact, 3 would only strictly be entailed by

- 2a. I condemn every person's inclination to cruelty.

But this commitment (unlike the more abstract and, as it were, Platonic 2) forces the issue of whether an individual person can be said to have the inclination to cruelty if the cruel wishes remain unacted on because of a countervailing restraint. In the light of this reflection, it looks as if 3 and *a fortiori* 4 should be clearly counterintuitive. (Compare: Nazi party membership is deplorable (because it leads significantly often to crimes against humanity); Oskar Schindler's membership in a Nazi party is deplorable; Schindler is to be deplored for his membership in a Nazi party.)

Nevertheless it turns out that 3 and 4 are not clearly counterintuitive, not even in the cases of Schindler and the mushroom, for the following reason. Even if we do not doubt A's sincerity in renouncing cruel acts, we may, recognizing A's imperfect self-knowledge and our imperfect knowledge of her, acknowledge the possibility that, in spite of appearances, her confession of inclination is a preliminary step towards the performance of cruel acts, or a symptom that these are impending. If we could be absolutely certain that the wishes would remain unacted on forever, they would seem to be the deadest of dead letters and to call for no comment at all—but in a human world of intellectual and moral fallibility this absolute certainty is never available, which is why Nagel's view that even the unacted-on inclination is condemned looks plausible. Similarly, we nervously abstain from eating the mushroom from the poisonous species even if a professor of toxicology reassures us that this particular one is safe, and we wish Schindler had not been (briefly) a party member because the blood runs cold at the thought of the crimes he might subsequently have been led to commit as a result of his membership, even if we know he didn't commit them. (Invoke possible worlds here, if it helps to make sense of the intuition.) In other words, we are tempted to endorse 3 and 4 in our intuitive judgments because the reservation "though unacted on" is as intuitively unstable as the police chief's promise not to shoot any additional prisoners. Condemnations 5 and 6, based on acts, leak into 3 and 4 and infect them (psychologically, not logically) with the image of acts. And condemnations 5 and 6 come into play in our intuitions in the first place because it is grossly anomalous (psychologically, not logically) to think of an inclination to a certain kind of act without thinking of some of the acts to which it is an inclination.

We can proceed, then, if I am right, on the basis that it is acts

that are the object of our first-line moral judgments. And it is both rational and conventional to judge acts at the moment of action, and in the light of all relevant circumstances, including the knowledge then available to the agent. I do not mean to deny that in practice moral condemnation or, more rarely, commendation may be founded on unreasonably high (or more rarely, low) expectations of an agent's decision making in the light of his epistemic opportunities or other circumstances. The possibility of such unfairness is itself a reflection of human fallibility, moral as well as intellectual. But it can in principle be recognized as unreasonable by reasonable people: there is a conventional discourse through which initial judgment is queried, corrected, and stabilized at least to a limited extent.

The examples of "moral luck" brought forward by Nagel and Williams lose their credibility if we bear in mind when analyzing them that agents are expected as a condition of moral commendation to take reasonable steps to deploy knowledge, including knowledge of their own and others' fallibility. This condition underlies the morally charged idea of *negligence*, exemplified by Nagel's case of the baby in the bathwater. Here the agent is in both cases (that in which the baby drowns, and that in which it survives) equally guilty of negligence, that is of taking insufficient care, by means of bathtime safety regimes, to guard against his fallible memory or powers of concentration on individual occasions. Certainly we know that we have "done something awful," as Nagel puts it, in the first case, but it is not so clear that we have "merely" been careless in the second. What we have done in the second case—and the cold sweat that can be provoked even by the memory of such a slip testifies to this—is something (act) that might have led to something (consequence) awful. Putting it like that suggests that Nagel's formulation of what we know we have done in the first case is a compressed version of "something (act) that actually has led to something (consequence) awful." In other words, case one is negligence multiplied by disaster, while case two is negligence multiplied by nothing much. The *emotion* we experience in the first case is of an appropriately complex kind, formed both by horror at what has happened and by self-reproach. (An observer might naturally try to capture this with the statement: "He *reproached himself horribly*, you know, after he failed to prevent that baby from drowning." The adverbial formula underlines the

confusion and mutual intensification of the two emotions. (Contrast “He was horrified; and he reproached himself,” which suggests a more rationally analyzed and, just for that reason, in the circumstances, a less natural or at least less appealingly “human” response.) In the lucky case, a judgment of self-reproach is equally appropriate, but the accompanying emotion, unalloyed with distress at a terrible outcome, can scarcely be comparable in intensity, let alone persistence: even the most vivid imagining of hypothetical disaster does not, except in certain neurotics, keep coming back unbidden as a bad memory of actual disaster does. These psychological facts (if they are admitted to be such) merely show that our emotional life is less analytically structured and differentiated than our rational judgments; they do not show that we cannot sustain judgments of our own culpability that are founded on a coherent idea of responsibility for actions. As for the judgment of an observer, a sufficiently reflective person will in the unlucky case mentally separate the culpability from the bad luck, employing probably one of those common formulas (“There but for the grace of God go I”) that seem to have been designed for the purpose.⁸ If there is in practice a tendency for fellow citizens to judge the culpability more severely in the event of disaster than in the harmless case, supposing (a significantly unlikely supposition) the latter were to become known, it is in principle correctable by reflection, information, or persuasion. (I will suggest later why the law may be justified in punishing negligence more severely when it happens to have harmful consequences.)

The recognition of human fallibility as a condition of moral decision and moral judgment helps also to clarify another example, this time from Williams. A truck driver accidentally kills a child: despite being innocent even of negligence, he will feel worse than any spectator, and though people will rightly seek to move him away from this feeling, “it is important that this is seen as something that should need to be done, and indeed some doubt would be felt about a driver who too blandly or readily moved to that position.” This looks like a clear case of intuitions validating “moral luck”: the driver is *wholly* without culpability, the position in which he finds himself is *wholly* a result of misfortune, and yet we expect him to feel distinctively badly, not merely (as a spectator

⁸André makes a not dissimilar point (204).

would) out of grief and horror at the death of a child, but because, as Williams puts it, "there is something special about his relation to this happening, something which cannot merely be eliminated by the consideration that it was not his fault".⁹ Indeed there is: his epistemic relation to an event in which he is a participant. If the truck driver is morally mature, he will be conscious of fallibility, of the possibility of errors of reasoning, perception, and other functions in morally significant situations, and will regard the requirement of deployed knowledge as being at its most stringent in the most morally compelling cases, such as those life-and-death contingencies of which a truck driver can hardly fail to be aware. His reason and his senses may tell him that he is innocent of negligence, but we would think less of him (morally) if he had complete confidence in his reason and senses under such circumstances. We expect his awareness of fallibility to tell him that he may have been negligent in some respect not yet clear to him, that even if after reviewing his actions he believes himself innocent this may be wrong, and that his interest in escaping an accusation of negligence may be distorting his judgment.¹⁰ These epistemic considerations also explain why it is appropriate, as well as compassionate,

⁹Williams, *Moral Luck*, 28.

¹⁰What applies here to negligence in relation to actions also applies to negligence in relation to dispositions that might cause actions: the condition of fallibility requires that we scrutinize our dispositions for the potential to cause harmful acts and apply this scrutiny with a certain perseverance in order to discount a likely tendency to self-serving over-optimism. R. M. Adams ("Involuntary Sins," *Philosophical Review* 94 (1985): 3-31) goes so far as to maintain that we can properly be blamed for dispositions, feelings, attitudes, etc., even when these are involuntary. I agree only in so far as the involuntariness is the result of a failure to deploy knowledge, including self-knowledge, to an extent that is reasonable given the gravity of the relevant issues: a person who involuntarily shows contempt to another, for example, may have imperfect self-knowledge (he is trying hard to express kindness, but is naively imperceptive of his own just-below-the-surface animosity), or imperfect knowledge of externals (the callow youth does not realize that his turn of phrase or facial expression is a contemptuous one). In every such case of self-scrutiny, however, a point is reached at which the requirement of deployed knowledge cannot reasonably be pressed further by an observer. Adams's view that a person may in any case be blamed simply for having a bad disposition (on the apparent analogy of a pool being blamed for containing poisoned water) is less easy to accept. As Nagel observes (*Mortal Questions*, 25), saying that a person is bad seems to be different from saying that he is a bad thing, that the occurrence of his actions, or the state of affairs that he exists, is to be deplored.

for others to reassure him: they are telling him that in this case, fallible though he is, his reason and senses are to be trusted. Fellow spectators are reinforcing his judgment with theirs; non-spectators are reassuring him that his faculties have not degenerated and that he is not being too lenient with himself on the point of deployed knowledge.

Perhaps a truck driver who has been through the full sequence of morally relevant thoughts and emotions indicated above, and emerged, supposing this to be possible, with an absolutely clear conviction that he is free from negligence, still feels awful, simply because he has been “the instrument” of another’s death, because the body lay under “his” wheels. (At the risk of laboring the point, this feeling awful must categorically be distinguished from any lingering conscious or unconscious fears of guilt through negligence—I suspect there is, in fact, no psychological statute of limitations in such matters.) Can this unhappiness nevertheless be called a moral sentiment—if we bear in mind that the word ‘instrument’ is, very exceptionally among its applications to persons, to be understood in absolute and exclusive contradistinction to ‘agent’, without the faintest suggestion of complicity with the force wielding the instrument? Both intuitive considerations and semantic utility seem to argue against calling it moral. If the truck driver expresses an unhappiness explicitly unconnected with fears of negligence, his friends will simply sympathize—for who could possibly bear being part of an accident that ends with a killed child? But if he begins to express it in terms, or even to accompany its expression with body language, suggestive of its resembling other *moral* sentiments, who can doubt that they will assume that he is “blaming himself” as an agent (or negligent), and resume their efforts to dissuade him from doing so by pointing to the absence of agency on his part? Their dissuasion marks the inhospitability of ordinary moral custom and practice to a notion of moral blame that would erode the distinction between being an agent and being (in the noncomplicit sense) an instrument.¹¹ The semantic point is really

¹¹Holly Smith cites by way of analogy such examples as parents’ feeling guilt at giving birth to a genetically defective child, or spouses’ blaming their partners for dying and “abandoning” them (570). These analogies have force; but I have tried to show that we do not need to employ comparison to such wholly irrational types of quasi-moral “judgment” in order to account for the emotions of the truck driver or the drowned baby’s

the same point: the concept of moral responsibility loses its coherence if this residual unhappiness is listed among "moral" sentiments. Of course, Williams's point is precisely to question that coherence. But the truck driver example was supposed to provide intuitive evidence against it: if no such evidence will stand up, then a decision to call the residual unhappiness a "moral" sentiment would be no more than semantic mischief-making.

The "Gauguin" example is more complex, but ultimately yields to analysis along the same lines. At the point of moral decision Gauguin (as we will agree to call this hypothetical figure) is confronted with a genuine dilemma. If we continue to assume an altruistic morality, the dilemma in its most crudely stated form is that if he stays with his family he contributes to their happiness, whereas if he leaves them and goes to the South Seas he contributes to the happiness of numerous art lovers, but he cannot do both. (Williams does not, it is true, translate the claims of art for Gauguin into these altruistic or utilitarian terms, and some will find it implausible that an artist's attachment of importance to his own career could be motivated by moral rather than egotistical considerations, but we must make some such assumption in order to test the claim that moral justification is subject to luck in respect of outcomes.) There is a familiar problem here about weighing utilities or comparing duties of different types, but we need not concern ourselves with it: suffice to say that Gauguin has to resolve this problem as he thinks best. But for Gauguin simply to formulate the dilemma in this way is to do less than can reasonably be expected of any moral agent. He must also deploy knowledge, including knowledge of his own fallibility, and in view of the comparative gravity of the issue this requirement applies fairly stringently. The following assumptions are made in the crude formulation of the dilemma, and each of them Gauguin needs to consider carefully.

- (1) It is not possible to become a great painter while staying with his family.
- (2) If he stays with his family he will contribute to their happiness.

parent: these have a rational moral element however alloyed with other feelings.

- (3) If he leaves his family he will contribute to the happiness of numerous art lovers.

What poses the dilemma is assumption (1): it is highly speculative, and Gauguin must weigh the possibility that in a year's time or as a result of a strenuous effort to adopt a different perspective it might appear to be false. To the extent that it might be false there is a case for staying. If either (2) or (3) has a higher probability than the other, that provides a weighting in favor of the corresponding course of action, though this may not override other kinds of weighting related to degrees of unhappiness or deprivation, numbers of persons involved, and so on. Along with any other considerations relevant to these two assumptions, such as the possibility of his drowning en route to the South Seas or the possibility that his family secretly hate him, Gauguin must take into account the possibility that the renunciation of one course of action will diminish the probability of carrying out the other successfully: that his regret at an abandoned career will poison his behavior towards his family, or that the thought of his deserted family will spoil his creative development in the South Seas. And in weighing these latter possibilities he must ask himself as honestly as possible whether he has the temperament to surmount such futile regrets.

All this suggests that Gauguin has a difficult decision to make, unless one believes, as perhaps many will, that there is an overpowering argument on one side (such as a special obligation to close relatives). But if we grant for the sake of the argument that the decision is difficult, and that Gauguin's fundamental value, the well-being of others, is right, then he is morally in the clear provided that he deploys knowledge of the world and of himself to the best of his ability. The uncertain outcomes that, according to Williams, give rise to retrospective justification or unjustification of what Gauguin did are precisely what Gauguin must contemplate in advance, so far as he can, when he decides what he ought to do. Williams is dismissive of the idea that Gauguin might at the point of decision be said to require "reasonable conviction" that he has creative talent, asks ironically whether Gauguin should consult professors of art, and concludes that "the absurdity" of such questions "expresses an absurdity in the whole enterprise of trying to find a place for such cases within the rules."¹² But the absurdity

¹²Williams, *Moral Luck*, 24.

lies only in any supposition that this and the other questions Gauguin must ask himself are capable of being answered in a way that will deliver predictive certainty, which we are already agreed for the purposes of the argument they are not. The point of asking and seeking to answer them is not to achieve predictive certainty, with or without the aid of professors of art, but to escape a justified accusation of culpably inadequate pursuit and deployment of knowledge given the gravity of the case. Gauguin cannot know that he is to be a great painter, but equally he cannot be justified in dismissing the epistemic issue, however intractable, and simply taking a leap into moral space.¹³ If it strikes us as objectionable even to contemplate sacrificing the claims of close relatives to a project requiring such elusive epistemic support, that is itself an epistemically informed moral judgment.

Suppose now that Gauguin, having made the decision as best he can, leaves his family, that they die in misery, that he loses his talent. Does this show retrospectively that he was morally unjustified? Or suppose that his family quickly survive his departure, and he becomes a great painter. Does this confirm retrospectively that he was morally justified? In both cases the answer is no. In the former case, it is probably true to say that had he acted differently his own value (the well-being of others) would have been better served. But this is simply a case of fallible human foresight. It is a practical but not a moral failure: morality does not require us to display infallible foresight, which is beyond any human being, but to do our best to deploy knowledge and to apply this requirement to ourselves with especial stringency when grave issues are at stake, and Gauguin has done precisely that. For the same reason, in the latter case Gauguin has no reason to compliment himself on the vindication of his moral conduct, though he can, of course, in an entirely trivial sense, like a successful gambler, compliment himself on the acuteness of his foresight. To say this is not, as Williams might fear, to separate moral from rational justification. On the contrary, the act that must be justified, at the point when decision leads into action, requires rational justification as a subordinate

¹³It may said that we sometimes admire leaps into moral space, at least when the leaper leaps out of the agony of a moral dilemma, and not merely through nonchalance. But that reservation is, I think, a sufficient concession to my argument. It is also possible that the admiration is not moral admiration.

part of its moral justification, and morality retains its sovereignty: this contemporaneous moral “justification” of action must not be punningly confused with the retrospective “justification” provided by various consequences of the action’s happening to turn out in such or such a way: that is just gambler’s luck.

It might be protested that neither Gauguin himself nor an observer would be likely, in the tragic circumstances of the former case, to think or say that his decision had been morally justified. Gauguin would reproach himself and the observer would condemn, or at the very least expect Gauguin to reproach himself. But here, as with the truck driver, the requirement to recognize human fallibility must be taken into account. We have stipulated for the purpose of the example that Gauguin has in fact deployed knowledge to the best of his ability. But it is not for him to say that he has done so: indeed, like the truck driver, he will ask himself, if he is morally mature, whether he may not have been self-interestedly self-deceived in making the decision, or have taken the stringency requirement too lightly. It will be difficult, especially when one is relying on memory for an account of one’s decision making, to separate with certainty the recognition that one has misjudged intellectually from the suspicion that one has devoted less energy to judging than might reasonably have been expected by others, let alone by oneself. And in one respect Gauguin is likely to be worse off than the truck driver: the latter can perhaps be reassured by the spectators that he was not negligent—this may be a relatively simple matter of observation—but no observer is likely to be able to testify to Gauguin’s having done his best in making the decision. Gauguin is therefore likely both to feel worse in the long term than the truck driver and to be treated less sympathetically by others (though an observer who is morally mature will bear in mind her own limited access to knowledge about Gauguin’s decision making, her own propensity to error, self-deception, and so on, before pronouncing condemnation). Certainly this is bad luck, but it is bad luck in respect of the esteem of others, not in respect of moral culpability: the low esteem of others for Gauguin is *in principle* correctable (as would be the low esteem felt by someone for a truck driver who was simply known to have “killed a child”) but there happen to be severe practical obstacles in the way of correcting it. (Still, Gauguin might be able to explain to a reasonable person why he took the decision he did, and the latter would not

be guilty of endorsing a moral offense if she concluded that Gauguin had made a justified decision, on the basis of a right choice of value and a deployment of knowledge to the best of his ability.) The case proves, then, not that morality is subject to luck, but that moral choice is often very lonely, and one cannot necessarily comfort oneself with the thought that in the fullness of time one's action will be praised by other people. No doubt some people are unlucky to face more testing, more lonely, and more frequent moral choices than others, and to incur incorrect but practically uncorrectable blame more often than others. But that is not "*moral luck*," since a fully informed observer would be in a position to make a just evaluation taking account of these variations in the agent's moral opportunities and challenges.¹⁴ It is just bad luck (or good luck) for a person's happiness that certain facts that would affect his reputation if known remain unknown.¹⁵

Finally, we need to examine the significance for the argument of the conventions governing the assessment of responsibility and blameworthiness in law. Where such conventions fail to "correct" for the impact of luck, they cannot be assimilated to instances of contingent and correctable misjudgment by individuals: rather they represent, it may be assumed, a considered, and more or less collectively acceptable, view of the accountability of agents for a range of morally significant kinds of action. And clearly judgments in law are not, or not wholly, regulated according to a luck-free conception of moral responsibility. A person who avoids the commission of a crime solely by virtue of a blunder remains unpunished; severity of punishment often seems to reflect the gravity of the crime's consequences (which may vary quite unpredictably)

¹⁴Compare "value-added" notions in the evaluation of public or commercial services, where rival organizations are ranked according to the *increment* of "value" to the citizen or customer, rather than according to indicators of total "value" achieved, since the latter tend unfairly to favor organizations operating in advantageous circumstances (such as unusually affluent or healthy or well-educated communities).

¹⁵Richards draws a related and important distinction between what an agent deserves and the way in which, given our "epistemic shortcomings" we are justified in treating him (201-2). A jury may be right to acquit through insufficiency of evidence a person who is, as a matter of fact, guilty: but this does not diminish the culprit's (moral) guilt, and the possibility of correct acquittal in such circumstances does not complicate the concept of moral desert attached to actions (or, as Richards would have it, to enacted character).

rather than an appraisal of an intentional act in the light of the knowledge available to the agent; negligence leading to disaster is generally punished far more severely than equal negligence without consequences, if the latter is punished at all.

The obvious, though as it stands over-blunt, answer to this objection is that law is different from morality, and that, consequently, legal examples do not tell against an idea of moral responsibility. The relation between moral responsibility and legal accountability is actually, of course, a complex one. No system of law has ever, so far as I know, attempted comprehensively to reflect moral desert (an enterprise that, among other things, would arguably require a schedule of rewards as elaborate as that of punishments). If we suppose, not too implausibly, that law is essentially an instrument for maintaining social well-being, then legislators are bound to weigh a number of second-order objectives directed to that aim. One will be to maintain a sufficient, but not necessarily an exact, correspondence of law with morality. (In order to have confidence in it and to identify with it, members of the society will need to feel that their stronger moral sentiments are not offended by its laws, but there may not be the same necessity for them to feel that that all their moral judgments are given legal force; and indeed they may recognize the disadvantages of such an arrangement, not least because of their awareness of their own and others' intellectual and moral fallibility.) Another will be to deter citizens from antisocial acts, and it is true that this can be accomplished only imperfectly by confining punishment to those "unlucky" acts that are successfully performed and that obtrude themselves on the attention of the law enforcement agencies. It would be better, so far as the deterrent objective is concerned, to penalize the decision making itself, irrespective of the success or failure of the attempted act; and this would have the further advantage of bringing legal penalties into line with judgment on moral responsibility. But a range of different second-order objectives, also aimed at social well-being, condemn this policy. It would only be workable, if at all, at the cost of continuous intrusive monitoring of every citizen's mental life, and this would impede the pursuit of virtually every other social good, from political liberty to economic efficiency.

Moreover, the fact of human fallibility in self-knowledge and knowledge of others underlies an important further argument for the maintenance of differentiations of punishment such as that,

cited by Nagel, between attempted murder and murder, even though the moral culpability of the agent in the first case may often be as great as in the second. If murder is punished more severely than attempted murder, there is an incentive to fail, and we know that successful performance of an action often depends, in ways the agent may not be fully conscious of, on the wholeheartedness of his commitment to its success. The distinction drawn by the law means that moral responsibility is often inconsistently treated, but it also probably means that a few potential murderers "blunder" semi-intentionally because they are aware of the severer punishment that awaits success. Social well-being seems then to provide an argument for subordinating scrupulously just treatment of moral responsibility to the objective of maximizing the incentive to fail. A similar argument supports severer penalties for negligence when it leads to disaster: legislators cannot rule out the possibility that in some cases the apparent bad luck of the disastrously negligent will actually be a consequence of marginally greater negligence, and it is worth trying to deter this.¹⁶ It is arguable that a deterrent effect could be achieved even more satisfactorily, and just treatment of moral responsibility more nearly approached, by treating all attempts to injure or kill others as if they had been successful, and all cases of negligence as if they had had the worst possible consequences. (And where the public interest in deterrence is very strong, as in the case of murder or life-threatening negligence, there is in fact a tendency to approach this position.) It would not, however, really bring punishment into line with moral judgment, partly because it would ignore the possible marginal responsibility just discussed, and partly because (again as a result of human intellectual fallibility and limited knowledge) it would make for a very stark punitive inequity, especially apparent in the case of negligence, between incidents that come to light and those that do not. And there are powerful social arguments against a penal system so unflinching and invasive: citizens, at least of modern Western nation-states, seem in general to be willing to accept some loss of deterrence and of strict justice in favor of the social breathing

¹⁶In practice, of course, where it is abundantly clear that the outcome of an offence has been drastically influenced by luck, courts do often attempt to reflect this in the sentence, giving greater weight in such cases to just treatment of the offender than to the slight danger of weakening deterrence.

space afforded by judicial mechanisms that are activated in the last, rather than the first, resort—that is, when the consequences of acts are manifest. Moral responsibility, in short, is only one conception contributing to the institutional conception of legal accountability. If it is intelligible to a reflective citizen, as it is, that the latter will sometimes diverge from the former for political reasons that outweigh, but do not negate, optimally fair assessment of moral blameworthiness, there is no basis for supposing that the collective consent given to legal accountability undermines the coherence of the luck-free conception of moral responsibility.

I have argued that such intuitive and legal judgments as appear to substantiate the notion of “moral luck” are, on closer scrutiny, derived from, or in principle correctable in accordance with, an idea of moral responsibility which would be acceptable to any reasonable person as the foundation for considered judgments. A central feature of this idea is that moral agents are to be judged in the light of the knowledge available to them at the point of decision to act, given an acknowledged underlying condition of universal fallibility in the pursuit and deployment of knowledge. This requirement does the same work—but within a coherent view of moral responsibility—done by the ghostly hand of “moral luck”: that of relating the ambition of moral judgment to apply consistently to all human conduct to the reality of a human condition subject to accident and a human mind subject to error. Sometimes, as they are fallible, human beings, in overhasty judgments, forget or misapply the requirement. But any kind of complex judgment is fallible in some comparable way: considered corrections of such errors of moral judgment do not necessitate any disturbance of fundamental principles or outrage to primary intuitions.¹⁷

University of Central Lancashire, England

¹⁷I am grateful to the Leverhulme Trust for making possible the study leave during which the early drafts of this paper were written, and to a number of anonymous referees for invaluable advice and criticism.