

## *Punishment for Criminal Attempts: A Legal Perspective on the Problem of Moral Luck*

THOMAS BITTNER  
Wellesley College  
106 Central St.  
Wellesley, MA 02481  
USA

### **I The law of attempts**

In the criminal law, the law of attempts is of comparatively recent vintage.<sup>1</sup> It is part of an important contemporary legal trend towards early intervention in the criminal process. There are now a substantial number of crimes on the books that, like the crime of attempt, only require that the perpetrator start down the road to carrying out his criminal intentions and do not require him actually to have harmed (or, in some cases, even identified) his victim.<sup>2</sup> Besides the law of attempts, these

- 
- 1 Its first appearance in English law was probably *Rex v Scofield* (1784) Cald. 397, which upheld an indictment for attempting to burn down a house. The *general* rule that 'all such acts or attempts as tend to the prejudice of the community, are indictable' did not appear until *Rex v Higgins* (1801) 2 East 5.
  - 2 Interestingly, the concept of early intervention is given its most extreme expression in the Bible story (Deuteronomy 21:18-21) of the stubborn and rebellious son who is headed for a life of crime and who, therefore, ought to be put to death (by stoning). 'It is better for him to die now, while he is still innocent of bloodshed, than to die later when he is already guilty.' He is 'judged on account of his future' (Talmud Sanhedrin, Ch. 7, 72A).

new crimes include conspiracy and solicitation, forgery and counterfeiting, the possession offences (drugs, burglary tools, counterfeit money, automobile master keys, etc.), even corruption of youth (children). In all these cases, the law is stepping in to forestall harm, rather than waiting until a harm has already happened.<sup>3</sup>

What is the significance of criminalizing these types of actions? One possibility is to think of them as a broadening of the legal concept of harm. Traditional crimes, such as murder, arson, and assault, were legal responses to violations of interests in life, property, and bodily integrity. The new crimes, such as attempted arson, grow from the recognition that certain behaviours like attempt, conspiracy, or possession violate society's interest in security, and therefore harm society, even when they do not violate a personal or property interest.<sup>4</sup>

Another possibility is to think of the introduction of these new crimes as the law's shifting its focus away from the harm element of crime towards what is really essential to the idea of crime: a faulty mental state, such as criminal intent. The criminal law's main purpose, on this approach, is to make society safer by reducing danger.<sup>5</sup> A person's actual behaviour is relevant to this purpose only as evidence of a person's dangerousness to society; it is evidence, you might say, of his criminal will. This means that there isn't any need for the law to wait for the actor to harm his victim before charging him with a crime, because the harm isn't, strictly speaking, an element of the criminal offence. If possession of counterfeit money is sufficient evidence of his dangerousness, then the law is justified in stepping in. It doesn't need to wait until he actually passes the fake bills.

The first way of seeing the development of the early intervention crimes thinks of them as a *supplement* to the basic list of traditional crimes, deriving their legitimacy from those older crimes. So, attempted arson is a crime, because arson is a crime, and, in general, the second order interest protected by criminalizing the attempt is an interest in the security of the first order interest protected by the recognition of the

---

3 George P. Fletcher, *Basic Concepts of Criminal Law* (New York: Oxford UP 1998), Chapter 10

4 See Michael D. Bayles, 'Punishment for Attempts,' *Social Theory and Practice* 8 (1982) 19-29, 23 for an explanation of this way of characterizing the interests violated by criminal attempts (not endorsed by Bayles, himself).

5 This is the approach taken in the *Model Penal Code*, for example. American Law Institute, *Model Penal Code: Official Draft* (Philadelphia: American Law Institute 1962).

original crime.<sup>6</sup> In contrast, the second way sees this new category of crimes as a potential *replacement* for the basic list of traditional crimes. So, on this view, arson is a crime because attempted arson is a crime; in fact, arson, stripped of its inessentials, just *is* attempted arson, because the success or failure of the attempt is irrelevant to its criminality. It is part of the purpose of this paper to argue that this second way of conceiving of criminal attempts (and of conceiving of the trend towards early intervention in crime, in general) is a mistake.

## II Punishment for attempts

In Canada, and in many other jurisdictions, the crime of attempt is generally punished less severely than the corresponding completed crime.<sup>7</sup> This feature of the criminal code is widely regarded among ordinary citizens as settled and acceptable policy. Nevertheless, some philosophers and jurists,<sup>8</sup> in particular those who conceive of criminal attempts in the second way described above, have argued that this general policy is arbitrary and (morally) unjust and that, therefore, criminal attempts ought (legally) to be punished equally severely as are their corresponding crimes.

In my judgment, the argument offered by these reformers is not persuasive and we are doing the right thing to punish attempts less severely than we punish the corresponding crimes.<sup>9</sup> In the next section of this

---

6 There are a few odd cases in which a failed attempt is illegal while the successful act is not. For example, in some jurisdictions, although not in Canada, attempted suicide is a crime while suicide is not. Treason, too, seems to fit this pattern, because while attempted treason is a serious crime, successful treason is never prosecuted. I think it is clear that neither of these 'exceptions' poses a real problem for the claim at issue.

7 There are a few exceptions to this rule, such as attempting to kill the Queen. Canada Criminal Code, section 24, contains the definition of attempt; section 46 defines high treason.

8 Joel Feinberg, 'Equal Punishment for Failed Attempts,' *Arizona Law Review* 37 (1995) 117-33; Sanford H. Kadish, 'The Criminal Law and the Luck of the Draw,' *Journal of Criminal Law and Criminology* 84 (1994) 679-702; David Lewis, 'The Punishment that Leaves Something to Chance,' *Philosophy and Public Affairs* 18 (1989) 53-67; Andrew Ashworth, 'Criminal Attempts and the Role of Resulting Harm,' *Rutgers Law Review* 19 (1988) 725-72; and Hyman Gross, *A Theory of Criminal Justice* (New York: Oxford UP 1978), among others.

9 Lewis does not actually claim that we ought to change our current punishment

paper, I lay out the reformers' main argument for changing our current sentencing policies. Then in the following sections, I critique that argument, tracing its main weakness to the argument's source in an inadequate treatment of the so-called problem of moral luck. The reformers' argument ultimately relies on the same ideas about moral assessment that give rise to this interesting philosophical problem. Once the argument is laid out, it will be obvious that these ideas cannot be used to support the reformers' position, because if there really were a problem with moral luck, we would be no more justified in blaming people for their criminal attempts than we would be in blaming them for the harms they cause. One beneficial side effect of working through these considerations as they are raised by the reformers in the context of the criminal law is that our justification for punishing people for the harms they cause becomes clearer. We are, then, better able to defend our intuitions about moral praise and blame against the challenge raised by the problem of moral luck.

Yielding to the reformers' argument and equalizing the punishment for attempts and completed crimes would be a significant step towards replacing the traditional categories of crime with the new crimes of attempt, conspiracy, possession, and so forth. This change would fall right in line with the second way of conceiving of attempts, as outlined above. Indeed, Joel Feinberg (1995) explicitly advocates taking the further step of replacement in his version of the argument for equalizing the punishment. My critique of the reformist argument and its approach to the problem of moral luck amount, therefore, to an indirect defence of the first way of conceiving of the law of attempts and other similar laws; it sees these new laws as a supplement to the traditional categories of crime and not as a replacement for them.

### III The reformers' argument

The chain of reasoning that supports equal sentencing for criminal attempts often begins with discussion of an imaginary pair of crimes which are exactly alike except that in one case the criminal succeeds in harming his victim and in the other case luck steps in and pre-

---

guidelines. Rather, he argues, in the face of the facts, that we already do punish attempts equally severely as we punish the corresponding completed crimes. Lewis, 'The Punishment.' For pragmatic reasons, Kadish is uncertain about whether we ought to change our sentencing policies, although he agrees that they are irrational and unjust. Kadish, 'The Criminal Law.'

vents his victim from being harmed.<sup>10</sup> Lewis' story of Dee and Dum is typical:

Dee takes a shot at his enemy, and so does Dum. They both want to kill; they both try, and we may suppose they try equally hard. Both act out of malice, without any shred of justification or excuse. Both give us reason to fear that they might be ready to kill in the future. The only difference is that Dee hits and Dum misses.<sup>11</sup>

Lewis argues that Dum and Dee are equally wicked and equally dangerous to society, and so they are equally deserving of punishment. Since the difference between their two stories (namely the amount of harm they happened to cause) was merely a matter of luck, it is unjust to treat them differently. Whether we think the proper justification or function or aim of punishment is retribution or deterrence (or anything else) we ought, morally, to punish them equally severely.

In general, the idea is that if punishment is based on culpability and, as Fletcher says, 'the only fair basis for culpability is the actions under one's control — that is, what one can be sure of bringing about,'<sup>12</sup> then these two criminals ought to be punished with the same severity. Feinberg puts the point this way, 'When such momentous matters as a choice between the death penalty and a minor term of imprisonment ride in the balance, we want the decision to be as free of arbitrariness as possible. Reliance on good or bad luck, or assigning weight to factors beyond the foresight or control of the central players, introduces an element of arbitrariness into court proceedings.'<sup>13</sup>

Now, the story of Dum and Dee is intended to model the conceptual difference between attempted murder and murder or between any attempt and its corresponding completed crime. So, the intended implication of deciding that Dee and Dum deserve the same punishment is that all attempts ought to be punished equally severely as their corresponding completed crimes. All other factors that characterise crimes (past history, wholeheartedness of effort, lack of justification, or whatever) have been held constant between the two cases. Certainly, these other factors should influence the amount of punishment appropriate for a crime. We can still take into account the cruelty of the act, for ex-

---

10 Feinberg, Kadish, Lewis, Ashworth, and others all argue in this way. Feinberg, 'Equal Punishment,' Kadish, 'The Criminal Law,' Lewis, 'The Punishment,' Ashworth, 'Criminal Attempts.'

11 Lewis, 'The Punishment,' 53

12 Fletcher, *Basic Concepts*, 173

13 Feinberg, 'Equal Punishment,' 119

ample, to help determine how severely we punish. But, the implication is that being a *successful* attempt *per se* should not have any influence on the amount of punishment deserved, because that is the only difference between Dee and Dum and *they* deserve the same punishment.

Andrew Ashworth gives expression to the general claim that as between those who succeed in harming their victims and those who try but fail to harm them, there is no morally significant difference: 'The essence of a criminal attempt lies in the defendant's firm intention to commit the substantive offence, and the failure to achieve that aim is invariably attributable to factors such as luck, chance, and misjudgement which do not significantly diminish culpability.'<sup>14</sup>

If this is right, the reformers further claim, then the law of attempts ought to be changed. We ought to stop thinking of failed attempts as lesser crimes than their counterparts. Cases like Dee and Dum show us that failed attempts deserve the same amount of moral condemnation as successful ones. So, since legal punishment ought to reflect what is just and avoid any element of arbitrariness or luck, our policy should be to treat attempts just as we treat their counterparts.

To sum up the reformers' argument:

- a. Dee and Dum ought (morally) to be punished equally severely.
- b. So, all attempts ought (morally) to be punished equally severely as their corresponding completed crimes are.
- c. So, our legal practices should be changed in such a way that attempts and completed crimes are punished equally severely.

---

14 Ashworth, 'Criminal Attempts,' 733. It is worth noting that Ashworth distinguishes between incomplete and completed attempts (such as Dum's), and he specifically restricts his claims about reform to the latter. This restriction allows him to avoid one of the objections I will raise in the next section. It also makes his position a less radical one than Feinberg's, because it requires him to retain the crime of attempt, at least for incomplete attempts. When Dum follows his victim, draws his gun, and takes careful aim, he is already guilty of attempt. The police should not be made to wait until after he has pulled the trigger (and missed hitting his victim) in order to arrest him. Ashworth also assumes the burden of drawing the distinction between incomplete and completed attempts, one that seems clear in the Dee/Dum type cases but not for many others, such as burglary, larceny, sexual assault, and others discussed in the next section.

#### IV The inferences

In my judgment, all the claims that comprise this argument are false and both of the inferences are invalid. My focus in this paper is on challenging the moral claims (a and b), but it is worth pointing out that even if you agree with the moral judgment about Dee and Dum and you agree that their case generalizes to a correct moral claim about all criminal attempts, you might still resist the inference to the conclusion about legal policy because you think that there is an important gap or difference between morality and the criminal legal system. That is, you might reasonably believe that it is a mistake to try to make our legal policies always conform to what morality requires. While the criminal law should strive to treat everyone fairly, it also has other important claims on its attention. This is not an uncommon kind of view; many people believe, for example that the exercise of prosecutorial discretion, such as in a case in which a plea bargain is made, is an unjust but necessary part of the administration of justice, if only because no society has unlimited assets to fund criminal prosecution.

If we accept the claim that Dee and Dum ought (morally) to be punished with the same severity, are we thereby also obliged to accept the generalization of that moral judgement to all cases of criminal attempts? Are Dee and Dum an adequate representation of the *general* difference between attempts and successful crimes? It may seem that consideration of these two cases allows us to get at the essential (or conceptual) difference between attempts and successful crimes. By holding all other factors constant, it looks as if we have thereby eliminated from our consideration all that isn't relevant to our judgement of the proper moral response to a failed attempt *per se*.

In fact, though, these two cases are not a good model for the general difference between attempts and completed crimes, for three reasons. First, many attempts differ from their corresponding crimes in being different actions. In the case of Dee and Dum, the two perform the same action (shooting the gun) but it has a different result for each of them. This is what makes it so seemingly plausible to claim that, since they did the same thing (with the same intent, etc.), they deserve the same punishment for it. But, for many other kinds of attempts, there is no question that the actors deserve different punishment, because they do different things. Murder, for example, is a crime of harmful result, but in many other crimes, the harm (if there is one) is not a *result of the actions* of the criminal, but rather the harm is *in the actions* of the criminal. So, for example, Dee and Dum aren't a good model at all for burglary (breaking and entering), larceny (theft), sexual assault (rape), kidnapping, hostage taking, incest, piracy, bigamy, trespass, and so forth. In all these cases, the completed crime involves by definition a different

act than the attempted crime. For example, burglary requires that the criminal actually enter the house, while in attempted burglary the criminal doesn't enter the house (he isn't able to, he is interrupted as he is about to, he sees something that deters him just as he is preparing to, etc.). For another example, incest requires that sexual intercourse take place, while in the normal case of attempted incest, there is no sexual intercourse. There can be no *Dee* and *Dum* of attempted burglary or of any other offence outside the scope of the crimes of harmful result.

Moreover, even for crimes of harmful result, the law recognises many actions as criminal attempts that fall far short of 'complete' attempts like *Dum*'s. For example, the Model Penal Code requires only 'a substantial step' towards the commission of the crime, and this language has been widely adopted. The Russian legal system, in the case of a serious offence, requires only that the criminal prepare, in some way or other, to commit a crime. So, successful murderers have often simply *done* more than unsuccessful attempters, and that means that the difference between them is not like the difference between *Dum* and *Dee*.

Second, in some cases, the legal requirements for the completed crime are different from the requirements for the attempt. For example, the crime of money laundering requires that there be some amount of stolen money involved, while attempted money laundering doesn't have this requirement (the money provided in a police sting operation qualifies). So, it would not even be possible for many would-be money launderers (who are guilty of attempt) to succeed at their task, since the actions they took in attempting to launder money would not count as laundering money. For another example, attempted murder requires the intent to kill, while murder is possible with only the intent to do serious bodily harm. So, if a person had the intent to cause serious bodily harm, it is murder if he kills his victim, but another person who acts the same with the same intent but (by luck) fails to kill would not legally be guilty of attempted murder. Here again we cannot get the kind of parallel cases needed to make the reformers' argument work.

Third, the case of *Dee* and *Dum* doesn't accurately model the difference between attempts and successes, because often there is a *reason* that an attempt fails; it is rarely through luck alone that the victim escapes harm. *Dum* misses even though he does everything in his power to hit his target; in this respect, he is an exceptional case. Often, the reason that an attempt fails is something that is or could be within the knowledge or control of the criminal. So, the reformers' argument doesn't apply to a case in which a would-be murderer carelessly leaves the label on the bottle of poison he hands to his victim and the victim notices this and the attempt fails for this reason. It doesn't apply to a case in which a would-be murderer shoots at and misses his victim because he never bothered to practice using his gun. It doesn't apply to a case in which he

misses on his first try and decides that he has had enough, rather than getting a little closer and shooting again.

If we turn from murder to a much more common form of deliberate human killing, namely suicide, we see strong statistical evidence that the difference between attempts and completed killings is not, in general, a matter of luck. Those who attempt suicide are preponderantly female, while successful suicides are mostly male. If the difference between success and failure in attempted suicide were largely a matter of luck, we would see a similar gender proportion in both categories.

It may seem that I am missing the point of using *Dee* and *Dum* to think about criminal attempts. The idea behind this thought experiment is to isolate the key feature — the amount of harm caused — in order to see what effect it has on our judgments about the two cases. But, it would be a mistake to decide about the proper punishment for attempts on the basis of this kind of conceptual difference alone. The proper way to make the comparison between attempts and completed crimes is to take into consideration all the factors, both contingent and conceptual, that are stable elements of the two categories in question and of the difference between them. If the difference between attempts and successful crimes happens to correlate with some contingent features of crimes, then it is necessary and appropriate to take those contingent features into consideration when formulating policy.

The reformers would say that if we wish to think about a half-hearted attempt that fails (such as the man who didn't practice shooting his gun), we ought to compare it to an equally half-hearted attempt that, by luck, succeeds, because only that sort of comparison isolates the aspect of the act in which we are interested. This is wrong, because if we wish to think about attempts and completed crimes in general, we ought, instead, to use a typical attempt or an average attempt or a representative sample of all attempts and compare it to a typical or average or sample completed crime.

After all, how much could we learn, for example, about the differences between male and female chimpanzees in general if we insisted on making the comparison only by thinking about *Ree* and *Rum*, two hypothetical chimps who are exactly the same in every respect except for the fact that one is male and the other female? We could not learn, for example, that male chimpanzees are, in general, heavier than females. Similarly, by using *Dee* and *Dum*, we miss the fact that, in general, attempters are less skilful, less committed, and less persistent than those who succeed in harming their victims. So, we aren't able to figure out that, in general, they deserve a lesser degree of moral condemnation.

We can conclude that while our consideration of the cases of *Dee* and *Dum* manages to isolate, in some sense, the conceptual essence of the crime of attempt, this is not the right way to think about attempts

in general. Attempts are often very different actions than their corresponding completed crimes; they are often based on different legal requirements; and it is just a fact that generally they fail for a reason, not by luck alone. These legal and contingent factors matter for the question of the proper punishment for criminal attempt in general. The category of criminal attempt is better understood, not in its austere, purely conceptual dress, but as it really occurs.

Yet, even if this is true of attempts in general, it still seems we must concede that for some exceptional pairs of cases our current policy of regarding attempts as lesser crimes than their corresponding completed acts gives us the wrong result.<sup>15</sup> For example, Kadish (1994) discusses the case of *Commonwealth v. Malone*<sup>16</sup> in which two men 'decided to play a game of Russian roulette in which each took turns spinning the chamber of a revolver, with one round in it, and firing at the other.' By law, the winner and the loser of this game would receive very different punishments for what they did, even though the difference between them is just luck. Is this aspect of our legal system something that we can live with?

It is pretty clear that we must, if we value the rule of law. Central to the ideal of the rule of law is the commitment on the part of the government to promulgate general rules of conduct that it will apply consistently in every case. The requirements of generality and consistency are quite important here; a system that judges individuals on a case by case basis would not realise the ideal of the rule of law nearly so well as one that operates on the basis of general rules. Given this, we cannot reasonably hope to develop a system of legal categories that automatically yields the correct results for every imaginable set of circumstances. Our aim, instead, should be to put in place the best system we can and to

---

15 Dee/Dum type cases are exceptional, but they are not just a philosopher's invention. One instance of them is when the victim of an attack is sent to hospital and his attacker is apprehended. The attacker may be charged with murder only if the victim dies. But, the victim's survival seems to depend on factors (availability and skillfulness of medical treatment, health condition of the victim, etc.) that are outside of the knowledge and control of the attacker. The attack is finished, yet the legal classification into which it falls is not yet determined. Therefore, this determination depends on luck, and that is what is unfair, say the reformers. A similar situation arises in connection with the use of Victims Impact Statements in sentencing. See *Payne v. Tennessee* (1992) 11 S.Ct. 2597. Here, the severity of the punishment in a criminal case can be affected by the harm suffered by the *family* of the victim, which is, again, a factor that is likely to be outside the knowledge and control of the criminal.

16 *Commonwealth v. Malone* (Pa. 1946) 47 A.2<sup>nd</sup> 445, quoted in Kadish, 'The Criminal Law,' 682.

build into that system a limited allowance for the use of discretion by its administrators.

Consider the case of Gee and Gum:

Gee really wants to hurt his victim; he's been brooding over it all week as he plans and rehearses his attack. Gum is just feeling a little frisky today and is momentarily peeved at his victim, who is a person about whom he really has no serious feelings. Seizing the moment, Gee taps a key on his victim's laptop, permanently deleting the only existing copy of a file that represents a few months' worth of his victim's labour. At the same moment, Gum half-kiddingly takes a swing at his victim, who is on the rugby team, causing a bruise, which quickly gets lost among all of his other rugby bruises.

Surely, Gee deserves greater moral condemnation than Gum does: he had more wicked intentions and he caused more harm. But, legally, Gum has committed the more serious crime of assault, while Gee is guilty of only a mischief. If anyone were to suggest, in the spirit of Dee and Dum, that this odd pair of cases shows there to be something wrong with our legal categories of assault and mischief, it seems clear that this would be a mistake. From the fact that, if we try very hard, we can think of an exceptional pair of cases where our legal categories don't reflect the actual differences in the moral blameworthiness of the two criminals involved, we should not conclude that our system is in need of reform. Gee deserves a more severe punishment than Gum. But, even if it were impossible under our current criminal laws to punish Gee more severely than Gum, that would not be a good reason to reorder our understanding of the difference between assault and mischief, or our understanding of the difference between crimes against the person and crimes against property.

## **V The premise**

We have seen that even if the reformers are right that Dee and Dum deserve equally severe blame (and punishment) for their actions, we could not use that lesson to support their more general contention about the appropriate amount of punishment for criminal attempts generally. In my judgement, these considerations constitute a decisive refutation of the reformers' argument in favour of changing our current sentencing policies. However, I also want to challenge the reformers' claims that are specifically about the relative moral blameworthiness of Dee and Dum. This topic leads us into the territory of the philosophical problem of moral luck.

In saying that these two scoundrels deserve an equal amount of moral blame for what they have done, the reformers are aware that they go

against what almost everyone would say, prior to reflection about this particular situation. I have never encountered any advocate of capital punishment, for example, who would be willing to say that Dum, just as much as Dee, deserves to die for what he has done. And yet, on a general level, there also seems to be a good deal of plausibility to what the reformers say, because they argue for it on the basis of a certain general intuition about moral assessment that seems clearly to be true and to apply to this situation. This general intuition is also the basis of the argument for the most common version of the problem of moral luck.

In my view, both the reformers and those who raise and defend the problem of moral luck as a challenge to our ordinary practices of moral assessment give an erroneous characterisation of this intuition and, for that reason, apply it incorrectly to our current punishment guidelines for criminal attempts and to our practices of moral assessment generally. I will argue that, correctly characterised, the intuition does not give rise to the problem of moral luck and it is actually consistent with our practices of moral assessment, including our pre-reflective assessment of Dee and Dum.

My main task for the remainder of this paper, then, is to examine three possible characterisations (or interpretations) of this intuition and the ways in which these characterisations might lend support to the reformers' argument. In the course of this discussion, I will also look at the problem of moral luck and its advocates' roughly parallel attempts to characterise and appeal to the same general intuition in support of their claims. It will turn out that these two groups make a similar misinterpretation of the intuition, and that this error is easier to see in the general context of the problem of moral luck, because it is applied more consistently there than it is in the context of the problem of criminal attempts.

The intuition in question is that moral assessment should not be affected by luck. Indeed, Bernard Williams has said that 'The *point* of [our] conception of morality is, in part, to provide a shelter against luck.'<sup>17</sup> Alternatively, this principle might be called the 'control condition' on moral responsibility: a person should be blamed only for those things that were within his knowledge and under his control and not for those things that were beyond his control (which would make them due merely to luck).<sup>18</sup>

---

17 Bernard Williams, 'Postscript' in Daniel Statman, ed., *Moral Luck* (New York: SUNY Press 1993), 251-58, 251. Italics mine.

18 For the sake of simplicity in framing the control condition, I ignore the special case of negligence, in which a person can be blamed for harms he did not foresee and did not try to produce. See fn 31, below.

The connection between responsibility and control that is expressed by the control condition runs deep in our conception of moral assessment. If this condition were not in place, a person might turn out, through mere happenstance, to be a morally bad person and this seems unacceptable. Martha Nussbaum calls this worry 'moral vulnerability' and traces philosophy's efforts to deny or contain it back to classical Greece.<sup>19</sup> Judith Andre describes Kant's view of morality along these lines as well: for Kant, morality is 'the sphere of life in which, no matter what our circumstances, each of us can become worthy. Moreover, moral worth is the highest worth of all, and so there is a kind of ultimate justice in the world: each person is equally able to achieve that which matters most in life.'<sup>20</sup> We cannot allow luck to foreclose this opportunity.

In his classic article on moral luck, Thomas Nagel describes the luck-excluding principle, which he also attributes to Kant, in this way: 'Kant believed that good or bad luck should influence neither our moral judgement of a person and his actions, nor his moral assessment of himself.'<sup>21</sup> As Nagel says, 'prior to reflection it is intuitively plausible that people cannot be morally assessed for what is not their fault, or for what is due to factors beyond their control.'<sup>22</sup>

With this principle in place, the problem of moral luck is supposed to arise because our ordinary practices of moral assessment seem plainly to give moral significance to many things that are due merely to luck. Only a little reflection is required to show that everything (or nearly everything) a person can normally be said to have done is at least partly determined by factors beyond his knowledge and control. Every action and consequence of an action, from growing a vegetable to killing an enemy (or victim), depends to some degree on the cooperation of factors that the actor could not have ensured or prevented. Nevertheless, it is our normal practice to assign praise and blame for many of these actions and consequences. For this reason, the line of thinking goes, our normal practice regularly violates the control condition expressed in the Kantian principle. And, now that we are aware of this inconsistency between practice and principle, it seems that we should either give up

---

19 Martha Nussbaum, *The Fragility of Goodness* (Cambridge: Cambridge UP 1986)

20 Judith Andre, 'Nagel, Williams, and Moral Luck,' *Analysis* 43 (1983) 202-7, 202

21 Thomas Nagel, 'Moral Luck,' in *Mortal Questions* (Cambridge: Cambridge UP 1979), Ch. 3, 24

22 Thomas Nagel, 'Moral Luck,' in *Proceedings of the Aristotelian Society Supplementary Volume* 50 (1976) 137-51, 138, and Nagel, *Mortal Questions*, 25

the principle or change our normal practices of assigning moral responsibility.

The legal reformers don't follow out the inference to this radical a degree. Their version of the argument retraces it in the particular circumstances of criminal attempts, but it stops far short of claiming that no one is morally or legally responsible for anything he does. Still, one can see the similarity between the two lines of thought by comparing Nagel's formulation of the Kantian principle to the reformers' explanations (in section 3 above) of why Dee and Dum are equally morally responsible for what they have done.

The reformers also wish to restrict moral responsibility to those events that are within the knowledge and control of the agent and in this way eliminate the influence of luck. Their proposal is supposed to achieve this result by identifying a class of actions that are, in their view, completely within the criminal's knowledge and control and for which he can be held responsible without violation of the Kantian principle. Since a person's success at killing someone partly depends on luck, he can't be held responsible for his victim's death, but he can, according to the reformers, be held responsible for having tried (attempted) to kill him, because the act of trying *is* completely under his control. Understood in this way, the Kantian principle implies the reformers' main claim: there is no moral (or legal) difference between attempted murder and murder or between failed crimes and successful crimes, generally, because person should be blamed only for what he really did and not for the consequences of his actions.<sup>23</sup>

## VI Is it a *Kantian* principle?

Many people have followed Nagel in regarding Kant as the main source of the idea that the consequences of a person's actions never matter for his moral status, but in fact, Kant did not hold this belief. It is true that in *Foundations of the Metaphysics of Morals*, Kant claims that a good will is

---

23 Since my concern in this article is with the reform movement as a whole, I have taken as my target an argument that is roughly common among reformers I mention in fn 8, above, and have not dealt with some of the specifics or the subtleties of Andrew Ashworth's arguments for reform. See his 'Belief, Intent, and Criminal Liability,' in J. Eekelaar and J. Bell, eds., *Oxford Essays in Jurisprudence: Third Series* (New York: Oxford UP 1987); 'Criminal Attempts and the Role of Resulting Harm,' *Rutgers Law Journal* 19 (1988) 725; and 'Taking the Consequences,' in S. Shute, J. Gardner, and J. Horder, eds., *Action and Value in Criminal Law* (New York: Oxford UP 1993). I will take up Ashworth's approach to the question of punishment for criminal attempts in a separate article.

the only unqualifiedly good thing in the world. If you have a good will, you cannot be blamed if, by bad luck, you do not achieve your purpose or even if your actions result in harm to yourself or to others. As Kant famously says, the good will 'would shine like a jewel for itself' even if (as it happened) it 'were entirely lacking in the resources to carry out its aim.'<sup>24</sup> After quoting this passage, Nagel makes the assumption (without offering any support for this interpretation) that Kant 'would presumably have said the same about a bad will: whether it accomplishes its evil purposes is morally irrelevant.'<sup>25</sup> That would give us a symmetrical principle of moral assessment by which 'there cannot be moral risk'<sup>26</sup> for either people of good will or people of bad will.

This assumption is wrong, however. Kant doesn't believe that the two cases should be given the same treatment. Rather, he says that if you have a bad will, you *can* be blamed for the bad effects of its exercise, even if those bad effects are unintended and even if they are only fortuitous. In *Lectures on Ethics*, Kant says,

if we do just what we ought and no less, any untoward consequences are in no way our fault and cannot be debited to us; we cannot be held responsible for them. On the other hand, if we do either more or less than we need, the consequences can be imputed to us, either as merit or as demerit....if I owe money and do not pay my debt when I ought, then if my creditor goes bankrupt in consequence, it is my fault: I did less than I ought to have done, and blame can be laid at my door.<sup>27</sup>

If the will is bad, Kant goes so far as to lay blame for fortuitous, unintended harms even if they were unforeseeable. Discussing his notorious example of someone who lies to try to protect the life of an innocent person, but inadvertently harms him, Kant says, 'whoever tells a lie...must answer for the consequences, however unforeseeable they were, and pay the penalty for them.'<sup>28</sup> In telling a lie, a person intends to deceive another person — he fails to treat him with respect — and therefore has a bad will. Aiming at good consequences (the 'altruistic

---

24 Immanuel Kant, *Foundations of the Metaphysics of Morals* (Indianapolis: Bobbs Merrill 1969), 1st section, 3rd paragraph.

25 Nagel, *Mortal Questions*, 24

26 Ibid.

27 Immanuel Kant, 'Responsibility for Consequences of Actions,' in *Lectures on Ethics*, trans. Louis Infield (New York: Harper & Row 1963), 59

28 Immanuel Kant, 'On a Supposed Right to Lie from Altruistic Motives,' in *Critique of Practical Reason and Other Writings in Moral Philosophy*, trans. Lewis White Beck (Chicago: University of Chicago Press 1949), 348 (quoted by James Rachels, *The Elements of Moral Philosophy*, 4<sup>th</sup> ed. (New York: McGraw Hill 2003), 124-5)

motives' mentioned in Kant's title, in this case, protecting an innocent life) notoriously doesn't, for Kant, morally justify the intention to lie or make it an instance of a good will. So, it seems clear that Kant sees the good will as protected from blame for inadvertent bad consequences, but he doesn't grant the same protection to a bad will. Certainly, then, he would be willing to lay the blame on Dee for the intended and foreseeable death of his victim, even though we must concede that it is merely a contingent, and not a logical, consequence of his actions.

This interpretation of Kant's principle, which I will call 'the asymmetrical Kantian principle,' puts Kant's position more in line with the mainstream western moral tradition and with our contemporary practices of moral assessment, including punishment for attempts, than Nagel's interpretation makes it out to be. People might still disagree with Kant about whether lying to save an innocent life is justified, but the disagreement would then be about the question whether altruistic motives might constitute a good will, which Kant (perhaps in opposition to what most people would say) would deny; it would not be about whether a person can be held responsible for whatever bad results flow from a good will, which Kant (along with most everyone) would, on my reading, also deny.<sup>29</sup>

Kant is only one of many important figures in the tradition to support the asymmetrical view of moral responsibility. St. Thomas Aquinas expresses the same view in this way: '...this is why the law lays down that if a man engages in legitimate activities and uses due care, he is not guilty of any homicide that may ensue; if, on the other hand, he engages in illicit activities, or even fails to take due care in some legitimate enterprise, he is guilty of any homicide that may occur.'<sup>30</sup>

---

29 It is beyond the scope of the present article to continue this defense of the attribution of the asymmetrical principle of moral responsibility to Kant. I'll just say that it seems to me that much more could be said, including the claim that this kind of asymmetry is a basic theme of Kantian ethics. The good will is guided by reason and genuine necessity, while the bad will is guided by self-love and inclination, which are merely contingent. So, it is entirely appropriate for contingency to be taken into account in assessing only the bad will and not in assessing the good will.

30 St. Thomas Aquinas, *Summa Theologiae* IIaIIae, Q. 64, Art. 8. In mentioning the issue of 'due care,' St. Thomas is raising an important and difficult problem for any general analysis of the conditions of criminal and moral responsibility. Why do we hold a person responsible not only for the harms that he intentionally causes but also for those harms that result from his failing to take due care in 'some legitimate enterprise'? Clearly, a person who intends the harm he causes may be held responsible for it, but we also blame the person who did not aim at or even foresee the harm he caused, *but could have and should have foreseen and prevented it*. In other

The asymmetrical version of the principle provides protection from moral vulnerability just as much as its symmetrical version does, and it maintains the same equality of opportunity to achieve moral worth that Kant envisioned. In order for luck to add extra moral blame to a person's actions, it is first necessary that he open himself to this possibility by having a bad will. A good will provides complete moral protection, on Kant's view. On the other hand, in having a bad will the agent has no guarantee of any limit to how much he may be blamed for. Because a bad will is, by definition, implicated in every crime, the asymmetrical Kantian principle (unlike the symmetrical version) justifies the general practice of punishing successful crimes more severely than their corresponding failed attempts. Good people may be protected from moral responsibility for unlucky harms they cause, but bad people, such as criminals, are not.<sup>31</sup>

Apart from the law of attempt, the criminal law generally seems to conform to the asymmetrical view of responsibility. A person can be held legally responsible for the bad consequences of his bad will even when those consequences are the result of luck. For example, a bank robber who makes a clean escape with his loot and doesn't physically harm anyone can be convicted of felony murder if the police officers

---

words, there are many circumstances in which a person is required to meet certain standards of due care and is held responsible for failing to meet those standards.

It is beyond the scope of this article to discuss this topic; my sketch of the conditions of moral responsibility offered below ('the weak symmetrical control principle' discussed in sections 7, 8, and 9) does not explicitly cover negligence (the failure to exercise due care). Control is a necessary condition for responsibility in the case of intentional harm. But, the case of negligence demonstrates that knowledge and control cannot be necessary conditions for moral responsibility in general, because in negligence, a person is held responsible even though he did not aim at or even foresee the harm he produced.

The main claims I make about foresight and control in the context of attributions of moral responsibility for intentional harms concern how we should understand these two concepts and, in particular, how they are related to luck. In negligence, responsibility does not depend on actual foresight and control but rather on the foresight and control that should have been exercised but was not. In my judgement, the claims I make about foresight and control will apply just as much to whatever additional principles are needed to account for responsibility under negligence as they do to the more basic case of responsibility for intended harms. See, also, fn 31.

- 31 My simple division of actors into those with good will and those with bad will leaves out the important category of those who are negligent or reckless (those who fail to exercise due care, as St. Thomas says) as well as those who may be held responsible under so-called 'bad Samaritan law' for harms they failed to prevent, but should have. Again, it is beyond the scope of this article to take on these complications.

who respond to the alarm shoot and kill one another by accident. On the other hand, a person is normally not held criminally responsible for any harm he may cause by chance when his mental state is faultless. There are many situations like this in the law, for example, when mistake is allowed as a defence in the case of actions that would otherwise be crimes. If, when leaving a restaurant, you make the innocent mistake of taking another person's umbrella, for example, you would not be prosecuted for theft.

The three main selling points of the symmetrical Kantian principle are that it offers us needed (and deserved) protection from moral vulnerability, that it seems to capture our intuitive sense that control (and knowledge) is required for moral responsibility, and that it seems to have the support of Kant and to be rooted in his overall approach to morality. We have now seen that Kant and the mainstream tradition don't hold the symmetrical view and that we can get whatever protection from moral vulnerability we need and deserve from the principle that they did actually hold: the asymmetrical principle. The Kantian vision of ultimate justice (or safety) does not require a symmetrical treatment of moral responsibility for consequences. It is consistent with both versions of the principle.

Unlike the symmetrical principle, the asymmetrical principle justifies our intuitions in cases such as Dee and Dum; it supports punishing Dee more severely than Dum. Unfortunately, the asymmetrical version of the principle has significant problems of its own. For one thing, outside the friendly confines of Kantian moral theory, it seems somewhat arbitrary to assign moral significance to the bad consequences of a bad will and not to do so for the bad consequences of a good one. A further troublesome objection to the asymmetrical version of the principle is that although it yields answers that fit with our intuitions in some of the cases we have been explicitly considering, there are other cases that it gets wrong. Indeed, in so far as our current practices of assigning criminal responsibility follow the asymmetrical principle, they also err, I claim. This problem with the asymmetrical principle stems from those features it shares with the symmetrical principle, in particular the fact that both maintain a sharp distinction between control and luck. I turn now to the question of the relationship between control and luck, and in section 10 I return to these 'problematic' cases<sup>32</sup> and end with my own suggestions for (mild) reform of the criminal law.

---

32 Joel Feinberg, 'Problematic Responsibility in Law and Morals,' *Philosophical Review* 71 (1962) 340-51

## VII Control and luck in moral assessment

The question of the need for control and the relevance of luck in moral assessment is an especially difficult one, because we get no clear guidance from our untutored feelings and attitudes towards the issue. On the one hand, in ordinary speech we distinguish between the act of murdering someone and the act of trying (but failing) to murder him. And, we certainly feel that there is an important difference between the two acts. But, these ways of talking may not reflect a difference that is even intended to be morally significant and these feelings may not be moral assessments. After all, a person may have bad feelings for another person simply because the other person harmed him in some way, even if he is convinced that the other person is not morally to blame for the harm he happened to cause.

On the other hand, it does seem plausible to many people to say that the moral assessment of a person's action should not depend on something that was outside of that person's control. Adam Smith, as Paul Russell reminds us, claims that our sentiments on this question are in conflict with one another. On a general or theoretical level, we may feel that only a person's intentions can be relevant for moral assessment.<sup>33</sup> But, Smith also asserts that on a case-by-case level, we have moral feelings that recognise a difference between the blameworthiness of Dee and Dum.

Asking ourselves to consider other, similar cases doesn't seem to help, either. For example:

Fee takes a shot at his enemy, Fie. He wants to kill Fie out of malice, without any shred of justification or excuse. At the moment he fires his gun, Fum (Fee's beloved brother) innocently and unforeseeably jumps in front of Fie, takes the bullet, and dies.

Is Fee just as morally blameworthy for having killed Fum by mistake as he would have been if he had successfully killed Fie as he intended? I think reasonable people's judgements may differ about this case. Fee's intentions are the same as Dee's, his physical movements are the same as Dee's, and he does cause a death by what he does in pursuing those intentions, so the causal condition and the harm seem to be the same. Still, as I watch Fee break down in anguish over what he has done, I

---

33 Smith says that this feeling is 'acknowledged by all the world and there is not a dissenting voice among all mankind.' Adam Smith, *A Theory of the Moral Sentiments* (1759), quoted in Paul Russell, 'Smith on Moral Sentiment and Moral Luck,' *History of Philosophy Quarterly* 16 (1999) 37-58.

think I am probably feeling that there is a significant moral difference between his case and Dee's. If my moral assessment should only take his intentions into account or even if I include his actions, then I should feel the same about him (and Dum) as I do about Dee, but I do not.

The symmetrical Kantian principle tries to make some headway in this difficult area by insisting, in the face of our conflicting intuitions in cases, that we should never ignore the control condition. Adopting the symmetrical principle means that we must make some changes in our practice of assigning moral responsibility, either the more moderate ones advocated by the reformers or the more radical ones that Nagel suggests. On the other hand, the asymmetrical principle can be seen as a kind of compromise between the control condition and our usual practices. It allows us to maintain a requirement of control in some cases and relinquish it in others. So, it sees luck as sometimes playing a role in moral responsibility and at other times not. Unfortunately, without the accompanying Kantian distinction between the necessity of the sphere of moral action and the contingency of the sphere of immoral action, this compromise is arbitrary.

The two versions of the Kantian principle we have been examining share the same basic conception of control and its relationship to luck. For both of them, control is an all-or-nothing condition, and control is regarded as lost with respect to something (S) whenever factors that are themselves beyond the knowledge and control of the agent have any influence at all on the occurrence or nature of S. In order for an action or one of its consequences to be suitable for moral assessment, according to the two strong versions of the principle we have examined, its occurrence or nature must be entirely uninfluenced by factors such as these.

This is quite an extreme view of what is necessary for a person to be in control of something (and therefore possibly to blame for it), because it assumes a sharp distinction between the agent and the world he inhabits that is reminiscent of metaphysical dualism in the philosophy of mind.<sup>34</sup> The model for this kind of view is a private sphere within which a person has 'genuine' control and luck plays no role at all in what happens. Outside the private sphere, a person may or may not be able to effect his designs, but we cannot say that he really is in control of what happens.

---

34 G. Ryle, *The Concept of Mind* (New York: Barnes and Noble 1949) is still the best critique of Cartesian dualism and what Ryle calls 'the two lives (or two worlds) legend.' The reformers are committed to something like the two lives legend, though they are perhaps not setting the dividing line exactly where Descartes would have had it.

An alternative to this kind of view would be to model the field without a sharp boundary around the agent and to see control and luck as mixed over the entire field with control normally more predominant in the area around the agent and luck more predominant in the areas away from him. Instead of thinking of luck as something sharply distinct from knowledge and control, this alternative conception sees them as intertwined in nearly every aspect of a person's life.

Anthony Duff also characterises the reformers' view as grounded in a dualistic account of action.<sup>35</sup> The view he criticizes says that since luck infects all the consequences of one's actions, and even bodily movements can always be subject to the influence of luck, the only thing a person can be said to do, strictly speaking, is to try. Human action, this view holds, consists of trying. So, a person can properly be held responsible only for what he tries to do. Trying is a purely mental act, on this view, giving us a dualistic account of action and moral responsibility.

According to Ashworth, the reformers' position can be formulated in a way that avoids this criticism, if a trying is conceived of as a physical event, such as pulling a trigger or shooting a gun.<sup>36</sup> Certainly this way of analysing our sample case comports with our natural description of what happened: in shooting his gun, Dum tried (but failed) to kill his victim. That was all that Dum was able to do in that situation, try. But, the reform position requires that we say more: it requires us to say that all *Dee* did either, strictly speaking, was pull the trigger or shoot the gun; all he did, strictly speaking, was try. This claim depends on our being able to draw a line between a person's actions and the consequences of his actions. And we can only do this if we are willing to say that it is only the shooting that was sufficiently within Dum's control and the killing was not. 'Each of them tried to kill, and it is that trying alone which is sufficiently within their control, whereas what happens in the physical world thereafter may be affected by other forces and circumstances.'<sup>37</sup>

However, this further claim doesn't comport at all with our natural description of what happened. When *Dee* pulls out his gun and deliberately shoots and kills his victim, it would be very odd to say that it was a matter of luck that he succeeded. This could be said only by someone who was already wedded to the idea that the bare possibility that something unforeseeable might intervene is enough to say that

---

35 R.A. Duff, *Intention, Agency and Criminal Liability* (New York: Oxford UP 1990), 188

36 Ashworth, 'Taking the Consequences,' 109

37 *Ibid*, 108-9. See also Ashworth, 'Belief, Intent, and Criminal Liability.'

even success is a matter of luck. Rolling a three in a game of dice is lucky; catching a taxi quickly on a rainy day is lucky. But, doing your laundry or opening a window is not lucky and neither is killing someone when you shoot them with a gun. This is true even though it might happen, through luck, that you try but fail to kill him.

It was Dum's bad luck (and his victim's good luck) that Dum's attempt failed. This doesn't mean that it was a matter of luck that Dee's attempt was successful. Dum could not, perhaps, have foreseen that he would fail, if the reason his intended victim survived was a real fluke (say, if the bullet was stopped by a lucky medallion hanging around his neck), but this doesn't mean that Dee could not have foreseen the death of his victim.

Ashworth asks 'whether it is possible to make an intelligible separation between certain "intrinsic" elements of behaviour or dispositions which are not subject to luck, or at least not subject to outcome-luck, and other "non-intrinsic" elements which are subject to outcome luck.'<sup>38</sup> His answer is to isolate trying, again understood not as a merely mental act, from the success or failure of results. For Dee and Dum, trying means shooting and success means killing. But, this model of a sharp border ('an intelligible separation') won't work, even for Dee and Dum, because luck touches the two of them in very different ways and they exercise different amounts of control and foresight in their respective situations.

So, the main idea I would like to propose is this: when a person does something, the occurrence of a certain result can be dependent on luck to some degree while at the same time being within his knowledge and control. Both the reformers' version and Kant's own version of the principle fail to recognise this possibility, because they embody a mistaken view of luck as something that stands essentially in contrast with agency (and therefore with responsibility). Both principles are based on the assumption that agents are, in an important way, distinct from the world they live in, not situated or embedded in it. According to this approach, only events that occur in the world outside the agent can be lucky or unlucky and only events internal to the agent can be under his genuine control, an expression of his free agency, and something for which he may be held responsible.

Both versions of the principle think of luck as the same thing as contingency and of control as the same thing as necessity. Indeed, Williams sometimes expresses the worry about moral vulnerability as a fear of the effects of *luck* on one's moral status and sometimes expresses it as

---

38 Ashworth, 'Taking the Consequences,' 108

a fear of *contingency*: '[our] conception of morality [is a] realm of value (indeed, of supreme value) that is defended against contingency.'<sup>39</sup> But, luck is not the same thing as contingency, just as control is not the same thing as necessity. Our basic intuition, which is correct, is that control is a requirement for moral responsibility. We go wrong when we conceive of control as sharply distinct from the rest of the world. We are well on the way to a better characterisation of the control condition, if we realise that control is a contingent circumstance, it is never completely free of luck's influence, and it is always a matter of degree.<sup>40</sup>

Understanding control as a necessary condition for moral responsibility but also as something always mixed with luck to some degree gives us a third possible version of the control principle: 'the weak symmetrical Kantian principle.' This characterisation of the control condition has some major advantages over the other two. A weak symmetrical principle offers the possibility of a better compromise between the pervasiveness of luck and the reality of moral responsibility than the strong asymmetrical principle is able to achieve. Under either principle, we get to blame the bad person for the foreseeable bad consequences of his actions (Dee gets the severe punishment he deserves), and we get to absolve the good person of blame for causing the unforeseeable bad consequences of his innocent actions (a person who innocently pushes the button to call the elevator and, because of faulty wiring, accidentally electrocutes its passengers is not to blame for what he has done). But, the strong asymmetrical principle also requires us to blame the bad man for the unforeseeable bad consequences of his bad actions (in the felony murder case, the robbers take the blame for the policeman's death even though they were not even in the vicinity when it happened) and it doesn't allow us to blame the good person for the foreseeable bad consequences of his permissible actions (the person in Kant's case who refuses to lie and tells the murderer exactly where his victim is hiding is absolved of responsibility for what he has done). If these latter two cases are regarded as unfortunate implications of the asymmetrical principle, and I think they should be, we would be better served to try another type of compromise.

---

39 Williams, 'Postscript,' 251

40 See R.A. Duff, 'Acting, Trying, and Criminal Liability,' in *Action and Value in Criminal Law*, ed. S. Shute, J. Gardner, and J. Horder (New York: Oxford UP 1993) 75-106, especially 101-2 and 106, fn 109, for a similar position on these questions. Winch argues that we should 'abandon the idea that criminal liability should depend on acts which are 'ours' in the 'metaphysically radical sense' citing the discussion in P. Winch, 'Trying,' in *Ethics and Action* (London: Routledge 1972). Duff's discussion is in the context of a general discussion of the nature of action.

The weak principle provides protection from moral vulnerability, just as the two strong principles do. Unlike either of them, it isn't based on an incorrect view of the relationship between luck (or contingency) and control (or agency). Unlike the asymmetrical principle, it doesn't draw an ad hoc or arbitrary distinction between the way the results of actions matter for the moral assessment of good people and the way they matter for the moral assessment of bad people. Like the asymmetrical principle (and unlike the strong symmetrical principle), it would support our current policy of punishing failed criminal attempts less severely than we punish completed crimes.

### **VIII Control and luck (again)**

Events and circumstances that are said to occur by luck are among those that are to some degree unpredictable or inexplicable or those that cannot be reliably produced or prevented. Though lucky and unlucky events or circumstances are very important to us in that they improve or worsen our lives, they also seem to be just those events for which a person should not be held responsible either because he didn't know something (as when you step outside and luckily catch a passing taxi) or because he couldn't control something (as when you roll the dice in a craps game and unluckily get a three). If something is foreseeable or predictable, then we are reluctant to call it luck. It is not luck if a person can know or guarantee that something will happen by his will alone.

If the key element of this model of luck is the border that marks the outer limit of the agent, then luck is contrasted with knowledge and control, which can occur only inside the agent. Within this limit, the agent knows everything and is in genuine control; outside it, the world is in control and the agent has no genuine knowledge. He can formulate all the 'best laid' plans he may wish, but to achieve his goals he must always hope that the world cooperates. In the model, the limit marks a qualitative difference between what goes on inside and out, between control and lack of control, between knowledge and guesswork. According to the strong symmetrical Kantian principle, the limit of a person's knowledge and control defines the limit of what he can appropriately be blamed or praised for. That's why this conception of luck must involve a contrast with agency across a boundary within which agency can supposedly work in isolation from luck.

Laid out in this way, I think this model of luck and agency is an extremely unnatural and unattractive one. It appeals to a picture of the self as standing outside of the natural causal nexus, interacting with the natural world, but also possessed of a private inner realm of knowledge

and control. Bernard Williams refers to this concept as 'pure agency' and denies that agency can be purified in this way.<sup>41</sup> Though he also gives voice to the opposing view, Nagel sometimes speaks as if it were possible to divide an action into the part caused by luck and 'that fraction of it which is under one's control.'<sup>42</sup> He tells us that 'from the inside we have a rough idea of the boundary between what is us and what is not, what we do and what happens to us...'<sup>43</sup>

Margaret Urban Walker argues that this model, which sharply distinguishes luck from agency, is a bad one. As Walker says, 'luck in the ordinary sense is not a whit incompatible with familiar, in some cases even criterial, indices of a responsible agent at its work.'<sup>44</sup> 'The notions of morality and luck are...deeply inter-related; for these notions are full members of a cluster of mutually sustaining concepts — the agency-matrix — which sort, and so stand (or fall) together.'<sup>45</sup> Far from being wholly distinct from agency, luck really only makes sense in the context of agency. In order to be lucky or unlucky, an event must have some kind of effect on a person's life. Not just any uncontrollable, unforeseeable event counts as lucky or unlucky. Luck must be someone's luck; it must either serve or frustrate someone's purpose or plan or interest. In this way, luck is 'wholly compatible with the exercise of rational choice, deliberation, wisdom, prudence...' etc.<sup>46</sup> (the lucky taxi-catcher was trying to get somewhere in a timely manner, the unlucky dice player was trying to win some money, and so forth).

Even more fundamentally, luck only makes sense in the context of agency, because luck's core feature of being in tension with knowledge, control, predictability, foreseeability, etc. means that it is defined in terms of 'agency-matrix' concepts. Once we dispense with the notion of agents with limitations on their abilities, it no longer makes sense to say of any event that it was a lucky or an unlucky occurrence. The distinction has no application.

Another way to demonstrate the essential interrelatedness of luck and agency is to try (and fail) actually to locate the boundary between

---

41 Bernard Williams, 'Moral Luck,' in *Moral Luck* (Cambridge: Cambridge UP 1983), 29-30

42 Nagel, *Proceedings of the Aristotelian Society*, 140

43 *Ibid*, 149

44 Margaret Urban Coyne [Walker], 'Moral Luck?' *The Journal of Value Inquiry* 19 (1985) 322

45 *Ibid*, 319

46 *Ibid*, 322

the inner world of agency and responsibility and the outer world of contingency and luck. In an early article,<sup>47</sup> Feinberg convincingly shows that luck is always involved with agency, no matter how much we try to remove and insulate agency from it. No matter how inwardly we try to locate the line to demarcate the area within which the agent is in complete control, luck and factors outside the agent's control penetrate the barrier and exert their influence. Once we see that luck and agency cannot be separated, that they stand or fall together, we can see that neither the strong symmetrical Kantian principle nor Kant's own strong asymmetrical principle is true.

In drawing the boundary of agency narrower and narrower, we aim to isolate those actions for which a person can properly be held morally responsible. The reformers want to draw this line just outside the attempt (successful or not) to commit a crime. They do this, because they believe that though the occurrence of the harm, itself, is a matter of luck, the actions that constitute the attempt are not. It turns out, as Feinberg shows us, that having tried to narrow the limit of pure agency to this point, we are still not completely free from luck's effects. Instead, we are only a little less subject to luck from our place within this boundary than we were outside it. This is still true even if we draw the boundary just outside of the agent's choices or even just outside of his intentions. We can never completely insulate those parts of ourselves that we care most about from contingency no matter how far away from the world and into the subject we withdraw.<sup>48</sup> There is no place to draw the line

---

47 Feinberg, 'Problematic Responsibility'

48 For an opposing view, see Thaddeus Metz, 'Arbitrariness, Justice, and Respect,' *Social Theory and Practice* 26 (2000) 25-45. Metz argues that we can draw the line at 'our essential nature as persons' where 'personhood, or rational will, is the ability to think and act according to deliberation' (28). Metz is engaged in the project of elaborating and defending an argument against libertarianism that turns on Rawls' notion of 'the arbitrariness of fortune.' See John Rawls, *A Theory of Justice* (Cambridge, MA: Harvard UP 1971). See also Robert C. Coburn, 'Distributive Justice and 'the Arbitrariness of Fortune',' *Philosophical Inquiry* 2 (1980) 441-57. This line of argument contends that it would be unjust to distribute the important goods and advantages of a society to its members in a way that makes the size of their shares depend in large part on their natural talents and abilities and their initial social position, because these natural and social endowments reflect merely the arbitrariness of fortune. I am sceptical of the prospects for success of this especially radical attempt to specify some sort of agent who is completely free of the effects of luck, because it leaves us with a very thin notion of the self as a pure rational creature. I doubt the claim that the ability to use deliberation in the service of thought and action is somehow essential to a person's nature in a way that his other talents and abilities are not. In any case, it is not at all clear that this view can provide any support to the specific kind of reforms I am opposing in this paper.

required by the Kantian model between responsible agency and the world, because agency is inextricably situated within the world.

Feinberg asks us to imagine a case in which one man (Hotspur) impulsively slaps another man (Haemo) during an argument they are having. Although Hotspur doesn't know it, Haemo is a haemophiliac and, as things turn out, the slap opens a small cut in Haemo's mouth and, since the bleeding cannot be controlled, he dies. Hotspur is found guilty of manslaughter. The fact that Haemo was a haemophiliac was merely a matter of bad luck, because there was no way that Hotspur could reasonably have known or become aware of that fact.

If we were to apply the strong symmetrical Kantian principle to this case we would say that Hotspur is not responsible for Haemo's death, because he could not properly be blamed for something that is outside his knowledge or control and that is, therefore, a matter of luck. Instead, he should be held responsible only for what he intended and chose to do, which was to slap Haemo in the face. This is where the reformers would have us draw the line. Hotspur is morally responsible for slapping Haemo, because that is what he did, strictly speaking; he is not responsible for Haemo's death, because that was a result of luck. Even though Hotspur did actually cause Haemo's death, it could have happened that he was not a haemophiliac. It could have happened that the slap did not cause the small cut; it could have happened that his life was saved by some new blood clotting drug. So, it was a matter of luck that Haemo died, because it was something that did occur, but might not have.

But, haven't we drawn the line too far out? Is the slap itself something that properly belongs within the limit of the agent's knowledge and control? Couldn't luck have stepped in to prevent the slap too, even in the context of Hotspur's intention and choice to slap him, in the same way that, as we have just seen, luck could have stepped in to prevent the death? It surely seems as if something might have happened to prevent the slap from occurring even as Hotspur is deciding to do it (a sudden muscle spasm, someone's stepping between them, an earthquake, etc.) and so it seems that the actual success of the slap must be taken to be a matter of luck, too, just as much as the death is. In order to avoid violating the principle, it seems we must draw the limit even more inwardly than the reformers would have us draw it. It could have happened that Haemo ducked and the slap never occurred. So, the slap is just as much a matter of luck as Haemo's death was.

Even if the slap, itself, belongs outside the boundary we are trying to draw, can't we still draw the line just inside the slap? Aren't the entirely internal events of formulating the intention to slap him and making the choice to slap him completely within the limits of knowledge and control of the agent and so insulated from the effects of luck? No,

even this extreme retreat into the agent is not enough to keep luck out. Consider, Feinberg says, some other possibilities. Couldn't it have happened that at the moment the anger is beginning to arise in Hotspur, a speck of dust causes a sneezing fit and prevents the intention to slap Haemo from occurring? Couldn't it have happened that just as Hotspur is about to make his decision to slap Haemo, he is distracted by a loud noise and by the time the noise subsides, his blood has cooled and he never chooses to slap him?

In the event, Hotspur did produce the intention and choice to slap Haemo, but it was partly luck that he ever formulated that intention and made that choice. So, it would also violate the strong symmetrical Kantian principle to hold him responsible for them. And, no matter how inwardly we try to draw the limit to what an agent can be held responsible for, we find that luck and contingency play a role there, too. Just as for any criminal act of harmful result we can always imagine an otherwise equivalent criminal attempt that, by luck, fails to cause harm, for every inner state, for every fully formed deliberation, intention, and choice to commit a crime, we can also imagine luck stepping in and preventing it, as well.

Conceiving of luck as what lies outside the agent encourages us to limit his moral responsibility to only those events that occur within that boundary and are therefore protected from luck's influence. This policy is our initial response to the realization that the effects our actions have out there in the world are always dependent to some degree on luck. So, we draw a boundary around ourselves and say that we are responsible only for those things we actually do and not for the effects our actions have on the world. When we realise that luck still plays a part in what we do within that boundary, we may reconsider and draw the boundary around what we try (or choose or intend) to do. But these turn out to be subject to luck, as well. There is no good place to draw this boundary; we ought to stop trying to draw it and concede that every aspect of our lives is touched by luck. Only then can we begin on the project of constructing a concept of ourselves as free agents who are genuinely situated in the natural world.

The application of these considerations to our original problem of punishing criminal attempts should, by now, be clear. We have been told that Dum and Dee should be punished equally severely for what they have done and that the reason for this is that it is the only way to make our practice of moral and legal assessment conform to the strong symmetrical Kantian principle. But, it now turns out, if this principle is true, then the same reasoning that shows that Dee deserves no additional blame or punishment for the consequences of his actions (the harms he caused) would show that Dum deserves no blame or punishment at all for his attempt. In both cases, factors outside the knowledge

and control of the agent contributed to whatever degree of success the agent enjoyed, whether it is in killing or in merely trying to kill.

## IX Applying the weak principle to criminal attempts

Once we realise that luck permeates every part of our lives — that we are limited contingent beings, in other words — we must adopt a principle of moral assessment that is consistent with that fact. How are we now to understand the difference between what we can be held responsible for and what we cannot be held responsible for? What are the relevant criteria to determine the amount of control necessary for moral responsibility? So far, all we know is that ‘perfect’ control is too much to ask for, since we don’t have perfect control over anything (not even ‘mental’ acts, such as deciding).

In trying to sketch the details of our third principle, we need to give content to a distinction between those consequences of a person’s actions that are within his knowledge and control and those that are not. Just as a person may be said to *know* about something even if he must concede that he might be wrong about it; a person may be said to *control* something even if factors beyond his control could step in and foil his aims. We should realise now that our ordinary concepts of knowledge and control can accommodate these kinds of possibilities.

Dee shot and killed his victim. Can he be held responsible not only for the shooting but for the death as well? He surely can, since the death was clearly what he aimed at and foresaw. He knew that the death was the likely outcome of his actions, and so in a perfectly ordinary sense of the word he *controlled* that outcome of his actions. It was something he brought about. In that sense, it satisfies a suitably weakened control requirement for moral responsibility. Our new principle will allow, therefore, that he is responsible for the death.

Dum missed. His victim did not die. What Dum controlled, as it turned out, were his movements; he shot intentionally and tried, but failed, to kill. We can hold him responsible for endangering his victim in trying to kill him, but not for harming his victim, since his victim was not harmed. This calls for a lesser degree of moral blame than Dee’s case does, because what he did was dangerous, but not harmful. A parallel case would be a person who had every good reason to believe something, but by some odd fluke, what he believed turned out to be false. If what he believed had been true, his belief would have counted as knowledge, in spite of the possibility of its being false. Since it was false, it was not knowledge.

Our weak principle is a symmetrical principle; that is, it treats the good will in the same way as it treats the bad will. In both cases, the

foreseen and controlled results of a person's actions are the only results for which he can properly be held responsible.<sup>49</sup> If a person, in trying to do something good, happens to hurt or to help someone in an unforeseeable or unlikely way, then it would not be appropriate to blame or praise him for it. Along the same lines, if a person who is trying to do something bad happens to hurt (or happens to help) someone in an unforeseeable or unlikely way, he is not to be held responsible for that, either.

We are still being judged by standards that are fitted to our all-too-human capacities. We still have some protection from moral vulnerability insofar as luck alone cannot take us from good intentions to moral failure. If something bad happens and we had good reason to think that it would, then we can properly be held to account, if we were in control of its occurrence. What we have given up is something that was unattainable in any case: complete control over and perfect knowledge of a realm within ourselves.

The strong asymmetrical principle makes those of bad will responsible for all the bad effects of their actions, whether those effects were within the knowledge and control of the actor or not. It lets those of good will off the hook for the bad effects of their actions whether or not those effects were foreseen (or even intended). In contrast, the weak symmetrical principle treats both good and bad will the same. It attributes responsibility for the foreseen or intended bad effects of good and bad will and it refrains from attributing responsibility for bad effects that were outside the knowledge and control of the agent.

It seems plausible to me that there will be cases in which a person does not deserve blame for the bad effects of his actions, precisely when he did not intend and could not have foreseen those effects. For example I think the liar in Kant's example, *even if we grant that his lie was a wrongful action*, doesn't deserve blame for the death of the friend he was trying to save. It really was just bad luck that the murderer found him just where the liar said he was. Similarly, Hotspur doesn't deserve blame for the death of Haemo, because there was no way for Hotspur to know that his slap would kill Haemo.<sup>50</sup> On the other hand, Dee does

---

49 Again, I here ignore the special case of negligence.

50 It would not be unreasonable to decide that Hotspur does deserve blame for the death of Haemo. If he does, it is because we can see a connection between what he foresaw or intended in slapping Haemo and what resulted (Haemo's death). If there really were no connection, if the case were like the elevator electrocution case mentioned in sections VII and IX and Haemo's slap were no more than a necessary condition of Haemo's death, we would not blame him for the death. I concede that reasonable people can disagree on whether that connection occurs in

deserve blame for the death of his victim, because he foresaw and intended and controlled his death.

Clearly, we would not want to attribute any unforeseeable and uncontrollable bad effects to an actor that was doing something morally required or morally permissible. When a person is 'minding his own business' exercising normal carefulness and thoughtfulness and through a fluke of luck happens to injure someone, he should not be blamed for it. The accidental elevator electrocution is our example of this kind of case.

It is a little harder to come up with cases in which an action that is permissible results in a bad effect that *is* attributable to the actor, but I think it does happen. The reason it may seem not to happen is that when you think of an example of a putatively permissible action that has a controllable or foreseeable bad result, the bad result tends to taint the putative permissibility of the action. That is, our utilitarian intuitions tend to make us see any action that has controlled bad consequences as being, by that fact alone, an impermissible act. In addition, whenever there is a foreseeably bad result, the action that led to it can be redescribed in such a way that under that new description it is an impermissible act.

Nevertheless, here is a possible example. When General Motors closed down their plants in Flint, Michigan in the 1980s (replacing them with plants in Mexico where labour costs were lower), they fired tens of thousands of auto workers and turned Flint into an economically depressed area, causing significant hardship. It was morally permissible for them to close the plants, since it was their property to do with what they wished. They broke no statutes or contracts they had made in laying off the workers. Nevertheless, I think it is arguable that GM was to blame for the economic hardship that resulted from their actions, since they could have foreseen that their actions would have just this effect.<sup>51</sup>

The weak symmetrical principle protects us from moral vulnerability by interpreting knowledge and control in the same sense we find used in ordinary people's judgments about responsibility. Real luck, the unforeseeable and uncontrollable, can't have a serious effect on a person's moral status. However, factors beyond a person's knowledge and con-

---

this particular case. Clearly it does occur in Dee's case and clearly it does not in the elevator case.

51 Compare the exposition of Kant's views on moral responsibility and the interesting and subtle discussion of a whole series of sample cases in Thomas E. Hill, Jr., 'Kant on Responsibility for Consequences,' in *Respect, Pluralism, and Justice* (New York: Oxford UP 2000) 155-72.

trol can play a role even in those things that he does know about and control. Like the asymmetrical principle (and unlike the strong symmetrical principle), the weak symmetrical principle is able to provide this protection while granting moral importance to some of the consequences of our actions as indeed we should, since they are just as much a part of what we do as are our actions themselves.

## X Upshot

For those who believe that our common practices of moral assessment and our current sentencing policies concerning criminal attempts are in good shape, the intuition that lies behind the Kantian control principle can give rise to cognitive dissonance. On its own, it will seem plausible, but when combined with the realisation that luck is pervasive in the human condition, it will seem suspect. There must be something wrong with it, if it is inconsistent with the basic outlines of our practices of moral assessment. I have offered an interpretation of the principle that doesn't challenge those practices.

An alternative way to resolve the tension is to assert that our common practices of moral assessment already conform to the (strong symmetrical version of the) principle. In a sense, I've partly employed that strategy in denying the inference from a to b in the reformers' argument (section IV, above). To say that most attempts are different acts than their successful counterparts is to say that for most cases, the way we currently assess moral responsibility for criminal attempts doesn't violate the strong symmetrical version of the Kantian principle, whether or not that principle is true.

This approach works well for the preponderance of cases, but there do seem to be some exceptional cases in which our practice still looks as if it violates the principle. If we punish Dum less severely than we punish Dee, then there is at least one case in which it looks very much as if the different treatment they receive from us is due simply to luck. I have argued that it only looks this way so long as we retain the sharp distinction between luck and control and that that there are conclusive reasons not to retain this sharp a distinction.

The reformers' argument asks us to regard Dee and Dum as a representative model of attempts generally and say, therefore, that all attempts differ from their successes merely by luck. This would mean that violations of the strong symmetrical Kantian principle are widespread, and therefore our practice is in need of reform. This is clearly not true, as I demonstrate above, because many failed attempts differ from their corresponding successes in other ways than by luck alone. So, even if the strong symmetrical Kantian principle were true and Dee and Dum

deserved the same punishment, that would not show very much about how we ought to treat attempts, generally.

The question remains what we should say about those exceptional cases, like Dee and Dum, in which the Kantian principle seems most clearly to apply. There, I think we are finally forced to examine the principle, itself, and whatever reasons we may have for assenting to it. Nearly everyone (except for Lewis) agrees that our current policies would administer a less severe punishment to Dum than it would to Dee, and nearly everyone (except the reformers) believes this is what the two of them morally deserve. If nothing else, this is a good reason to begin to suspect that any principle which condemns this practice is false.

The principle appealed to by those who would reform our practices has a good idea at its core: control *is* necessary for moral responsibility. Unfortunately, they take control to be incompatible with the bare possibility that a person's designs should be foiled. I have argued that this is not the case.<sup>52</sup>

*Received July 2005*

*Revised March 2006*

*Revised August 2006*

---

52 I thank two anonymous reviewers for this journal for their helpful comments on an earlier version of this essay.

