Moral Luck and Liability Lotteries

Fate and Fortune, two average drivers identical in every respect, went out on a ride taking the same route 5 minutes apart. They were both distracted by the same shoe advertisement and negligently took their eyes off the road. Whereas nothing happened when Fortune was driving, a slowdown in the traffic just when Fate was looking at the ad caused him to bump into Hurt, a motorcyclist who claims $5,000,000 in damages[1]. It's intuitively plausible that Hurt has a moral right to be compensated by Fate and that fate has a correlative moral duty to make reparations to Hurt, but that no such duty pertains to Fortune. Furthermore, the intuition that Fate has this duty to make reparations that are related to the amount of harm caused (let us call it the Moral Duty to make Reparations, or **MDR** in short) does not seem to depend on any particular legal order that validates this duty as legal; it seems to apply even in a 'state of nature'.

The problem is that many people also hold the view that luck may not affect our moral standing[2]. The thought that luck cannot affect our moral standing is described by many as a bedrock of our moral thinking and is famously ascribed to Kant[3]. Luck, in this context, consists in everything that lies beyond the agent's control, and although it is not quite clear what moral standing is, it is clearly related to things like blameworthiness and praiseworthiness. Let us call followers of this line of thought Adversaries of Moral Luck (or **AMLs** in short) and those who endorse the view that luck can affect one's moral

[1] The problem is taken from Jeremy Waldron, 'Moments of Carelessness and Massive Loss' in David G Owen (ed), *Philosophical Foundations of Tort Law* (Clarendon Paperbacks, Clarendon Press, Oxford 1995) (hereinafter: **Waldron (Moments)**).

[2] Here is a very partial list of AMLs: Joel Feinberg, HLA Hart, Judith Jarvis Thomson, Michael Zimmerman, Harry Frankfurt, Nicholas Rescher and others <<REFERENCES>>.

[3] See e.g. <<Enoch, Sher>>. For the orthodox view that Kant rejected moral luck see Nagel (Moral Luck) <<REFERENCE>>. For a more controversial line of thought that denies that Kant is opposed to moral luck see John Gardner <<REFERENCE>>.

standing Friends of Moral Luck (**FMLs**)[4]. The AML premise can then be described as follows: "One cannot be more or less blameworthy or praiseworthy because of things beyond one's control". Put differently: "If the only difference between agents A and B lies in things beyond their control, then A and B are equally blameworthy (praiseworthy)".

Where is the problem, you may ask. Well, *if* one accepts the AML premise *and* one also holds that the reason that Fate ought to compensate Hurt – the justification that underlies MDR – is that it is Fate's fault that caused the accident, then the problem is that to the extent that Fate is at fault, so is Fortune[5]. If the duty to make reparations is blameworthiness-related[6], then it should apply to Fate and Fortune equally[7]. At this point some may wish to invoke the concept of causation. The reason that Fate (but not Fortune) should compensate Hurt is that Fate *caused* the harm that Hurt suffered. But this seems to be begging the question, because that is exactly what the AML seems to deny. The AML denies that one can be more or less blameworthy because of things outside his control and the effects that his actions cause are clearly beyond the agent's control.

In this paper I will examine one solution that was suggested to the problem that faces the AML[8]. This is a solution offered by Jeremy Waldron that builds on an argument made to the same effect in the context of punishment by David Lewis. In the first section I will explain the solutions and how they aim to justify different punishment for successful and unsuccessful attempts and tort liability, respectively. In the second section I will argue

[4] Here is a very partial list of FMLs: Bernard Williams, Thomas Nagel, Martha Nussbaum, Tony Honoré, John Gardner, George Sher, Robert Adams and others <<REFERENCES>>.
[5] It may be worth noting that most non-consequentialist accounts of tort law rely on the premise that the tortfeasor ought to compensate the victim because he is at fault. See e.g. <<REFERNCES>>.
[6] In this paper I use fault, culpability and blameworthiness interchangeably.
[7] The same argument can obviously be made with regards to other blameworthiness-related institutions such as punishment of successful and unsuccessful attempts (see Feinberg <<REFERENCE>>) and other, non-institutionalized, forms of reaction such as attitudes like blame, guilt etc.
[8] I will, therefore, take for granted that the AML premise is correct. I will also assume, for reasons that will become clearer later on, a libertarian view on the problem of free-will. On the connection between the problem of moral luck and the problem of free-will see <<REFERENCE>>.

that the arguments way of justifying differential treatment of criminals and tortfeasors cuts under the main ideas that seem to underlie these institutions and that, at any rate, they only work if we presuppose some political order that validates these institutions renders them justifiable. The arguments will thus not work in the absence of such order, in a state of nature. Finally, I will argue in the third section that even if these problems could somehow be overcome, the account of Lewis and Waldron ultimately fails to solve the main problem that it set out to solve, that of moral luck.

1. The Tort Liability That Leaves Something to Chance

The solution suggested by Waldron is based on David Lewis' attempt to solve the problem of moral luck in the context of criminal attempts[9]. Under the AML premise, agents who try to commit a crime (say, murder) are equally blameworthy whether the attempt was successful or not. Hence, the unsuccessful criminal who tried to kill his victim is as blameworthy as the successful murderer, provided, of course, that the failure was due to factors beyond the agent's control (like a bird who unexpectedly flied in the course of the bullet, or that the victim surprisingly wore a bulletproof vest)[10]. How, then, can it be justified to impose different punishments on the complete murder and the unsuccessful attempt, as is the case in most common law jurisdictions? Here's one way of explaining our practices of punishing:

Suppose that every person who attempts a murder (whether successful or not) is to be punished by subjecting him to a punishment lottery. If he wins, he gets a lenient punishment. If he looses, he gets a severe one. Whether such system of punishment can be justified is arguable. It seems quite undesirable that some criminals will be more

[9] David Lewis, 'The Punishment That Leaves Something to Chance' (1989) 18 Philosophy and Public Affairs 53 (hereinafter: **Lewis (Punishment)**).
[10] <<Feinberg – REFERENCE>>.

severely punished than others just because of their bad luck. However, it is important to note that this is *not* a matter of moral luck but rather plain luck[11]. All the AML premise contends is that luck cannot affect one's blameworthiness; it does not deny that luck can affect one's well-being[12]. As far as the AML is concerned, the important thing is that the successful and the unsuccessful criminal are equally blameworthy. Because the successful and the unsuccessful criminal are equally blameworthy and their moral standing is identical, they ought to be subject to the same punishment and that is exactly what happens. The punishment that they are both subjected to does involve chance, but the chances are equal for both of them and the successful criminal has as much chance of getting the lenient punishment as his unsuccessful peer. Like other spheres of life, then, in punishment too, luck will affect one's well-being. However, it will not affect one's moral standing and hence will not be moral luck.

Even assuming that there is some way of justifying this weird system of punishment, it will still be very different from our system. In our system there is no draw, and successful criminals are simply punished more severely than unsuccessful ones. So a few modifications are needed. And here they are in brief:

First, the punishment imposed on all criminals is not a single-chance lottery, but rather finely attuned one, that takes into account their level of wrongdoing. So an attempter that subjects his victim to a 20% risk of death (because his hands are shaky, he uses a rusty rifle etc.) will be subjected to a 20% risk of severe punishment; one who subjects his victim to a 80% risk of death (firm hands, good rifle) will be subjected to a 80% risk of severe punishment and so forth.

[11] I develop the distinction between plain luck and moral luck in my thesis <<REFERENCE>>.
[12] Some AMLs (luck egalitarians) in fact make this further claim, but it is not an essential component of the AML view.

It may be argued that there is no reason to put the 20% risk imposer (let's call him Tom) under more favorable treatment (better chance of getting the lenient punishment) than the 80% risk imposer (Ernie). After all, Tom would have probably been happy to have firmer hands, more reliable rifle and better overall chances of killing his victim. Someone who sets out to kill someone, wants to have him dead rather than to expose him to some risk. Both Tom and Ernie are acting in the same manner – they are both doing their best to get their respective victims dead[13]. That one of them is more potent than the other is a matter of luck and cannot affect their respective degrees of culpability[14].

This argument invokes another kind of luck. So far we were concerned with resultant moral luck, luck in the effects of one's actions and how they turn out. The problem that the case of Tom and Ernie raises is that of circumstantial moral luck, luck in the circumstances in which one finds himself[15]. The problem can be described as follows[16]:

P1: If the only difference between Tom and Ernie is a matter of luck (i.e. something beyond their control), then Tom and Ernie are equally blameworthy (the AML premise).

P2: Ernie subjected his victim to a 80% risk of death.

P3: Tom too would have subjected his victim to a 80% risk of death if circumstances allowed (his hands were firmer, he had a better rifle etc.).

P4: The only difference between Ernie and Tom lies in circumstances that are beyond their control (from P2 and P3).

[13] Cf. Michael J Zimmerman, *An Essay on Moral Responsibility* (Rowman & Littlefield, Totowa NJ 1988) 49.
[14] In fact, if this argument is valid then both Tom and Ernie are as blameworthy as someone who subjects his victim to a certain (100% risk of) death. The argument that follows can be easily amended to accommodate that by replacing Ernie with another hypothetical character who does subject his victim to a 100% risk of death and by replacing the figure 80% by 100%.
[15] On the taxonomy of moral luck see Nagel (Moral Luck) <<REFERENCE>>.
[16] This is an adaptation of an argument raised by Zimmerman in <<REFERNCE – taking luck seriously>>

P5: Hence, Ernie and Tom are equally blameworthy.

Although this argument is only brought here to cast some doubt over one of the steps in Lewis' argument, I believe that its force is much greater. This argument, if valid, shows that the problem of moral luck is much greater than the problem that Lewis is trying to solve. If the argument holds, then the question at the beginning should wear different form altogether. Instead of asking why Fate is treated differently from Fortune, we should ask why Fate is treated differently from other drivers who (luckily) happened to take a different route this morning or from other drivers who were not distracted simply because they were lucky not to notice the advertisement (not to mention the pedestrians).

I cannot adequately deal with the problem of circumstantial luck in this framework[17]. What I will try to show here is that there is some plausibility to the thought that circumstantial luck does not pose a problem to one who holds a libertarian view of free-will. The argument can then be read as follows: Lewis' account works best to justify differential punishment if we hold a libertarian worldview[18]; I will show, however, that the argument ultimately fails even under the favorable assumption of libertarianism.

How, then, can a libertarian solve the problem of circumstantial moral luck[19]? The argument rests crucially on the premise that Tom would have subjected his victim to a 80% risk of death if circumstances allowed (P3). In fact, the premise as stated is inaccurate. It should instead state that Tom would have *freely* subjected his victim to a 80% risk etc[20]. This is a counterfactual that holds that Tom would *freely* do something given some hypothetical circumstances. Counterfactuals of this kind are commonly

[17] I tackle this problem in another paper <<REFEERENCE>>.
[18] The same goes for Waldron's account that aims to justify tort liability.
[19] The argument in the next two paragraphs draws on Adams and van Inwagen <<REFERENCES>>.
[20] If Tom would not have done so freely, he would, presumably, not be blameworthy at all, because blameworthiness presupposes freedom.

referred to counterfactuals of freedom and libertarians deny that they ever have truth values (or, alternatively, argue that they are always false). Clearly, if P3 (or any premise of a similar form) can never be true, the argument does not hold.

Why do libertarians hold that counterfactuals of freedom are never true? The reason is that libertarians hold that free actions are not deterministically caused[21]. An action is free, according to libertarians, only if, given the agent's intentions, desires, beliefs etc., there is more than one course of action available to the agent. Now if more than one course of action is available to the agent, what makes one course of action more 'true' than the other? What renders the premise that Tom would have freely chosen to subject his victim to a 80% risk etc. true and the opposite premise that Tom would have freely chosen *not to* subject his victim to a 80% risk etc. false[22]? Nothing, because there is nothing in Tom's intentions, desires, beliefs etc. that precludes him from freely deciding *not to* impose that risk on his victim[23]. Hence, because libertarians presuppose indeterminism in the process of free action, they seem immune to the problem that circumstantial moral luck poses in reliance of counterfactuals of freedom[24].

After this long diversion it is now time to return to Lewis' argument. As we may recall, Lewis suggested that criminals should be subjected to a lottery that is attuned to the level

[21] If they are caused at all.
[22] Some may think that if Tom *did*, in fact, subject his victim to a 20% risk of death, this shows that he would have, necessarily, subjected his victim to the greater 80% risk of death. A similar argument is made by Alvin Plantinga. Plantinga argues that if someone succumbs to a bribe offer of $10,000 then he would have also succumbed to a bribe offer of $30,000 and, hence, that at least some counterfactuals of freedom have truth values. The argument is refuted by van Inwagen <<REFERENCE>>.
[23] More technically we can say that no possible world in which Tom decides to impose the higher risk is closer to actuality than all possible worlds in which he chooses no to do so.
[24] Libertarians do, however, seem vulnerable to a variant of this problem. Although libertarians deny that counterfactuals of freedom have truth values, many of them hold that such counterfactuals have probabilistic values. Thus, even though a libertarian would deny that Tom would have freely imposed the higher risk, he may agree that he would have *probably* done so freely, and if that is true, then Tom is *probably* as guilty as Ernie. I deal with this variant in <<REFERENCE>>.

of risk that they actually imposed (and under the assumption of libertarianism, deny that there is such thing as the level of risk that they would have counterfactually imposed).

But that is still very different from our punishment system. Under our system there is no draw. So Lewis suggests a second change. We don't actually bother to organize a lottery. Instead, we use the results of the criminal's act as a substitute. If the attempt was successful we note to ourselves that the criminal has lost the draw and will be subject to the severe punishment and vice versa. This is a convenient arrangement because, presumably, the risk of death imposed by the criminal on his victim is exactly the same risk of severe punishment to which he should be subjected.

It is important to note that the criminal does not get punished more severely because his attempt was successful but rather because he was misfortunate in his draw. Indeed, there may be some confusion because the criminal act serves a dual purpose in the process of determination of the punishment. First it determines the criminal's level of culpability and the appropriate level of punishment that fits this level of culpability. This is the same whether the attempt is successful or not. Second, it serves as the determinant of the criminal's draw. Here success makes all the difference. If Tom kills his victim, the 20% risk he imposed on him materializes. That does not affect his culpability. Simultaneously, his punishment draw is performed subjecting him to 20% chance of severe punishment and he loses (the draw is the same apparently indeterminate process that was activated by his shot). The loss of the draw too does not affect his culpability but it does affect the actual severity of punishment that will be imposed on him.

This story, thus construed (and much abridged), is Lewis' attempt to justify our punishing practices. Waldron takes it a further step and applies it to our practice of imposing tort

liability[25]. Let's say that negligent drivers who are morally responsible for their negligent behavior (such as Fate and Fortune) should be subjected to a liability lottery. Given that Fate and Fortune are equally negligent, they should be treated equally. They both subjected the driver in front of them to a certain risk of injury and loss – let's say it's a 5% risk of $5M loss – and they should both be equally subjected to a similar risk of loss. Since there are probably different kinds of risks to which they expose their preceding respective drivers (e.g. 50% risk of $100 loss to property, 20% risk of $100K minor wounding etc.), they too should be subjected to several lotteries representing various levels of risks and correlating harms. Whether any risk materializes or not is determined, as Lewis suggests, by the actual act. Fortune's good fortune gets him off the hook. When he returns his gaze to the road it turns out that he won the draw (actually, all the draws) and won't suffer a loss. Fate is not so lucky. When he stares at the ad instead of looking at the road he is culpably negligent and subjected to a lottery in which he bears a risk of 5% of losing $5M (among other lotteries). Unfortunately for him he draws the wrong card (which is determined by the fact that he hits Hurt) and therefore suffers the loss.

Note that the number $5M represents both Hurt's loss and Fate's fine[26], but this is a mere contingency. Hurt isn't subject to a $5M fine *because* he caused $5M harm to Hurt. Rather, the same bad luck that affected Hurt (the slowing down of the traffic at the exact moment when Fate was watching the ad) also affected Fate and determined the results of his draw(s) and it's because of this bad luck that they both suffered their respective losses. We could have used a different algorithm to determine the results of Fate's draws and, had we done so, Fate may have been subjected to a different fine (or even no fine at all,

[25] Waldron (Moments) 401-405.
[26] I use the term fine even though fine usually represents some form of punishment. Indeed, under Waldron's scheme imposition of tort liability looks very similar to a punishment and, in fact, he sometimes calls it punishment.

like Fortune). Again, as in Lewis' explanation of punishment, Fate and Fortune get differential treatment because of plain luck, even though their moral standing is similar.

Waldron's solution seems to bypass the main problem nicely. It allows us to maintain that Fate and Fortune are equal in terms of moral standing and yet explains why we treat them so differently (because of plain luck). It thus aims to overcome the main obstacle – if Fate is no more at fault than Fortune, why is he treated in a harsher manner? The answer is: he is simply unlucky. The question that we should now ask is does this suffice to justify the vast difference between the treatments of the two. We will now turn to consider it.

2. Liability Lotteries – Are They Justifiable?

Both Lewis' and Waldron's solutions rest on the assumption that some independent justification can be found for establishing a system of penal/liability lottery, respectively. That a penal/liability lottery is not unfair and treats people who have identical moral standing alike, isn't enough. It might, for example, treat everyone unfairly in a similar manner. Or it could fail to manifest any worthwhile moral principle. Or it could simply be inferior to a system that accorded each person what he deserves, without invoking chance.

For that reason, Lewis says that a penal lottery might be justified if it "gets the case over and done with quickly"[27]. For example, if instead of building prisons we could subject criminals to various levels of risk of death penalty. He admits that this justification won't work if criminals are subject to varying periods of imprisonment.

Waldron, by contrast, offers a non-instrumental justification. According to Waldron, imposing a risk of loss on those who impose such risk through their negligence can be

[27] Lewis (Punishment) 59.

justified as a matter of *Lex Talionis* (hereinafter: **LT**) – an eye for an eye[28]. In a preceding article Waldron argues that LT is a "principle about what counts as an appropriate punishment" which is compatible with various theories of punishment[29]. He further argues that LT need not require an exact reenactment of the offense with the criminal as the victim, but rather imposition of punishment which is similar to the offense in some respects (the scope and extent of which, are open to a normative debate). For example, if A kills B on a Tuesday, it doesn't seem required by the principle of LT to execute him on a Tuesday. But the abstraction can go way further. Murderers, need not necessarily be subjected to capital punishment, according to Waldron's understanding of LT. They can instead be subjected to incarceration which will reenact several morally important features of the act of murder (disruption of life projects, grief-causing etc.)[30].

Precisely because LT is such a broad concept it cannot, in and of itself, justify the establishment of liability lotteries. Some further argument is needed to show that the most appropriate way to treat risk imposers is by exposing them to similar risks. It actually seems plausible to argue to the contrary, that since the morally relevant feature of risk imposing is that on occasion it results with causing harm, we should expose risk imposers to harm (rather than to a risk of harm). Waldron himself raises these possible objections and he replies that macabre appropriateness allows us to subject risk imposers to the exact level and structure of risk they impose on others[31]. But, again, why is that a morally relevant feature of their wrongful behavior? Isn't the salient wrong-making feature of negligent driving the harm that it causes rather than the risk that it imposes? If, indeed, LT requires that we choose a punishment that subjects the wrongdoer to the wrong-

[28] Waldron (Moments) 402-403. In fact, Waldron uses the doctrine of LT to justify Lewis' account of punishment, but I shall consider it as a justification for imposing a liability lottery on tortfeasors.
[29] Jeremy Waldron, 'Lex Talionis' (1992) 34 Arizona Law Review 25 (hereinafter: **Waldron (LT)**).
[30] Ibid 41.
[31] Waldron (Moments) 405-406.

making features of the wrongful act[32] it seems that we should subject tortfeasors to harm rather than to risk.

Waldron admits that it may be preferable to require *both* Fate and Fortune to compensate Hurt and his counterparts (rather than subjecting them both to a liability lottery) but before considering his suggestion, I would like to make another comment. Lewis' justification rests explicitly on institutional considerations (it is cheaper to have death penalty rather than imprisonment penalties). Waldron's suggestion rests on institutional considerations in a more implicit manner (it's more convenient to use the outcomes of the actual event to determine the results of the liability lottery). It should be borne in mind that the strength of these arguments is significantly diminished when one is considering the application of MDR in a 'state of nature' rather than the legal institution of tort law. Institutional considerations usually bear their full strength in a political setting where just institutions provide substantial goods and should generally be supported. In a 'state of nature', by contrast, the fact that the framework of a lottery is hard to ascertain seems to lead to the conclusion that the lottery system should be abandoned altogether. A lottery system seems to require some sort of political agreement which is lacking in a state of nature. In the absence of a political order Fate may well reply to Hurt "Look, it's true that I negligently caused you harm, but Fortune was just as negligent as myself. It may seem fair that negligent people should pay for their negligence, but I never gave my agreement to pay by means of lottery. In fact, had you asked me whether I am interested in such mechanism I would have probably said no; I would have said that I prefer to pay a small fine whenever I act negligently rather than a huge fine if harm ensues". This seems to pose a serious problem in the state of nature, but it may pose a problem in the political

[32] Waldron (LT) 33-35.

state as well. The arguments made by Fate seem sound and it seems likely that if we could choose through some political process whether to adopt a lottery system, we would have decided not to. For that reason Waldron comes up with another suggestion which we will now consider.

So far, we have ignored the question of why Hurt ought to be compensated. Liability lotteries can, in principle, be established even if no one is "Hurt". The proceedings of the lottery will simply go to the State's coffers or something like that. But let's suppose that there is some justification for compensating Hurt. The question remaining would be how the burden of compensating Hurt should be borne. If we phrase the question in this manner we see that liability lotteries provide an answer as to why it is fair that Fate bear the burden of compensation rather than both Fate and Fortune.

But if we do consider the question in this manner we must also consider whether it is better to establish liability lotteries instead of spreading the harm among all drivers (or all negligent drivers). As Waldron rightly argues, most drivers would prefer to spread the harm among all of them instead of taking the risk of bearing large portions of the harm caused as a result of negligent driving. In fact, most of them already do so by purchasing insurance policies[33]. Liability lotteries can be invoked to explain, then, why it is not unfair to impose liability on Fate, but they cannot explain why it is justified, all things considered, to do so.

Needless to say that liability allocation is even more remote from the intuitive thought that we started with, the intuition that supported MDR, than liability lotteries are. It also goes without saying that it seems hard to conceive that such scheme will bear any moral

[33] Waldron (Moments) 406-408.

weight in the state of nature. But apart from that, the problem with this line of argument is that it removes the problem of individual responsibility for outcomes but only at the expense of introducing the problem at the collective level. If the amount of harm that Fate caused cannot be an indication of his liability because it is subject to luck, neither can the total amount of harm created by all negligent drivers indicate the total liability because it, too, is a matter of luck. We can easily think of weeks when the issue of car accidents is mentioned in the media again and again and, by contrast, other weeks when all goes well and the negligent drivers and their victims get lucky. Whatever time-slice we choose for accumulating harm and allocating liability, the amount of harm caused will remain, fundamentally, a matter of luck[34]. Moreover, what reason is there for differentiating between negligent driving and other forms of negligent behavior, or even between negligent behavior and intentional harmful behavior? Presumably, whatever justifies compensating Hurt would justify rectifying any harm caused through wrongdoing and whatever justifies allocating liability among negligent drivers would also justify allocating liability among wrongdoers more broadly. And why should we confine the collective liability of wrongdoers to harm that they caused? In the absence of causation as a determinant of liability there seems to be no reason to allocate wrongfully-caused harms rather than harms more generally among wrongdoers.

The problems with Waldron's attempted justifications of liability lotteries demonstrate powerfully the Scylla and Charybdis of tort liability. In order not to violate the AML premise, the AML tort theorist must let go of causation in either its individual or collective form, but once causation is discharged nothing seems left to explain why the

[34] It is obvious that allocating liability among negligent drivers must be done within some time frame. There is no way of measuring the total amount of harm caused in the past and future unless it is delimited within a specified time frame.

reparations made by the wrongdoers are identical to the harm suffered by the victims of the wrongs. Liability lotteries that suggest the connection is only contingent provide an imaginative solution, but they seem to provide an ad-hoc solution that cannot be justified on its own grounds. Furthermore, as I will now turn to argue, even if liability lotteries could somehow be justified, they must ultimately be rejected by AMLs.

3. Liability Lotteries and the AML Worldview

The underlying basis of both Lewis' and Waldron's proposals is that the amount of risk imposed by wrongdoers can serve as a measurement of the level of culpability that their wrongdoing entails. To some extent that seems intuitively right. Grossly negligent driving is more culpable than marginally negligent driving because it bears greater risk of harm. Hence, both Lewis and Waldron maintain that we should expose the wrongdoer to the same amount of risk that he imposes on his victim – the more risky his behavior, the more culpable he is and, correspondingly, the greater risk of punishment/liability that he will incur. That seems sound, but how will the risk imposed by the wrongdoer be measured? Lewis makes several suggestions: (1) It could be the objective probability of the occurrence of harm; or (2) The reasonable degree of belief for a hypothetical observer who has all the possible relevant information regarding the situation; or (3) The reasonable degree of belief for a hypothetical observer who has all the information the wrongdoer actually has; or finally (4) The wrongdoer's actual degree of belief whether reasonable or not. It would be nice not to have to decide, says Lewis, and suggests that a reenactment of the wrongful act (or using the act itself instead of reenacting it) will impose the same level of risk on the wrongdoer irrespective of how we decide to measure it.

But it does not. Think of the following situation. Abe plans to kill Ike. For that purpose he buys a gun learns how to operate it and carefully plans how he will snare Ike. It so happens that the Big Brother Police learns about the plan and decides to let Abe carry out his plan (in order to frame him) while taking measures to protect Ike. Unbeknownst to Abe, a Big Brother agent enters his house and replaces the bullets in his gun with identically looking empty cartridges. In the next morning Abe carries on his plan. He meets Ike in a dark alley and shoots him at the head from range zero. To Abe's surprise the bullets don't even injure Ike and seconds later a police force shows up and apprehends him.

The objective probability of the occurrence of harm in this case was zero[35]. This is also the reasonable degree of belief for an observer who knows the relevant facts (most importantly, the replacement of the bullets). On the other hand, Abe believed that his chances of killing Ike were 100%[36] and so would believe every rational observer who knows the same facts as Abe and is unaware of the replacement of the bullets by the policeman. Now what's the relevant level of risk for determining Abe's culpability? It seems that the relevant level is 100% as in option 4 – Abe's actual degree of belief. Abe wanted to kill Ike, took the measures to do it and, as far as he knew, he was about to succeed in carrying out his plan. If his moral standing is determined by his response to moral reasons he seems to have openly rejected them in pulling the trigger[37]. He is no different, in that regard, from his counterpart Abe[1], whose gun is loaded with real bullets and whose shot kills Ike[1]. The penal lottery he should be facing is one in which he will (almost) certainly get a severe punishment. But the penal lottery he *actually* faces is one

[35] Or nearly zero. Let's leave supernatural occurrences outside the story.
[36] Or almost 100%.
[37] Cf. Zimmerman (Responsibility) 41-43.

in which he is bound with (almost) certainty to get the more lenient punishment. His *actual* chances of killing Ike (and thus losing the draw and getting a severe punishment) are zero. It turns out that the actual act does not provide an adequate substitute to an established penal lottery. Why is that?

The reason is that the wrongdoer's degree of culpability is not measured by the actual level of risk he imposes on his victim, but rather by the subjectively perceived level of risk that he imposes (Lewis' fourth option). An act is more culpable not when it imposes a greater risk of harm in the actual world but rather when it is done with the intention of imposing a greater risk of harm (whether or not it succeeds in carrying out that intention). When the result of the draw is determined by the outcomes of the act, it reflects the risk implicit in the act itself and not the risk that is intended by the agent. It subjects the wrongdoer to a level of risk that is proportional to his *doing* but not to the culpability implicit in his *wrongdoing*.

The problematic difference between the agent's actions and his culpability pertains not only in cases of intentional wrongdoing but also in cases of negligent behavior. Consider the case of a negligent driver Fate[1]. Fate[1]'s story is very similar to Fate's with one important difference. Seconds before his mind is distracted by the shoe ad, he passes by a signpost warning him of a slippery road ahead of him. Unfortunately, Fate[1] does not see the signpost through no fault on his side – there was a storm the night before (of whose existence Fate[1] was unaware) that caused the signpost to fall. We will never know whether or not Fate[1] would have abided by the signpost had he seen it[38], but since, in fact, he didn't see it, he doesn't slow down. Note that Fate[1] does not incur blame for not slowing down. Since he is ignorant of the bad conditions lying ahead, his decision to

[38] As you may recall, libertarians deny that counterfactuals of freedom can ever be true.

continue to drive normally and not to slow down[39] does not manifest a failure to respond properly to moral reasons. His ignorance exculpates him[40]. Seconds later Fate[1]'s mind is distracted by the ad and, at the same time, the road becomes slippery. When the traffic slows down in front of him and Fate[1] realizes this and tries to stop, it is too late. His inadvertence combined with the bad road conditions render it impossible for him to stop on time and he hits Hurt[1].

Fate[1] behaved wrongfully. He shouldn't have let himself be distracted by the ad. He should have also been driving more slowly due to the slipperiness of the road, but he cannot be blamed for that. What is Fate[1]'s degree of culpability? Let us look at the level of risk created by his behavior. Suppose it is 50% risk of causing harm[41]. When someone doesn't mind a slippery road he has 50% chance of causing harm. But not all of that risk is due to Fate[1]'s wrongful behavior. Some of it is due to facts regarding which Fate[1] isn't to blame. Had the road not been slippery, the level of risk imposed by Fate[1]'s negligence would have been 5%. If we were to establish a liability lottery for Fate[1] we would have designed it so that he will have a 5% chance of incurring liability. As things were, however, Fate[1] was subject to a liability lottery where his chances of incurring liability were 50%. Again, as in the case of intentional wrongdoing, the risk caused by the agent's actions is not necessarily a good measure of his culpability.

We can now generalize this point and conclude that penal/liability lotteries fail to assure that the wrongdoer gets his due. They rely upon the agent's actions and the risk of harm that they create as they manifest in the real world. But according to the AML worldview, culpability is not determined by how things are manifested in the real world but rather by

[39] Or, rather, the lack of a decision to slow down.
[40] On the exculpating effects of ignorance see Michael J Zimmerman, 'Moral Responsibility and Ignorance' (1997) 107 Ethics 410.
[41] For the sake of brevity I ignore the possibility of different levels of risk and different levels of harm.

how they are perceived in the agent's mind, since what happens in the real world – including what exactly is left to chance – is a matter of luck. The only thing that is under the agent's control is what goes on in his mind, and even his actions, as they are perceived from the outside, cannot adequately reflect this. If we hope a lottery to overcome the problem of moral luck it will have to be a lottery that is not based on the act and its outcomes but rather on the level of culpability which is entirely internal.

Conclusion