

Epistemological Contextualism and the Problem of Moral Luck

Berit Brogaard

Pacific Philosophical Quarterly 84 (2003), 371-83

Abstract: We have a strong intuition that a person's moral standing should not be affected by luck, but the fact is that we *do* blame a morally unfortunate person more than her fortunate counterpart. This is the problem of moral luck. I argue that the problem arises because account is not taken of the fact that the extension of the term 'blame' is contextually determined. Loosely speaking, the more likely an act is to have an undesirable consequence, the more its agent is to blame. But how likely a consequence is depends on which possibilities of harm we take to be relevant.

Introduction

You are driving down the street. You know perfectly well that you should have checked your brakes a long time ago. But you failed to do so. You had too many other things to think about. It's been raining. You are driving quite slowly, partly lost in your thoughts. Suddenly, there is a pedestrian on the street. You slam down the brakes, but they fail and you hit the pedestrian and kill him. Are you responsible for the death of the pedestrian? We would ordinarily say that you are—that you should be blamed more because you have killed a pedestrian than if you had merely been driving around with ineffective brakes. But can this be right? Can the mere fact that a pedestrian was crossing the street precisely then be enough to justify our assessing you differently from the point of view of blame?

The question raised here touches on the problem of moral luck, a problem which, as it is discussed in contemporary literature, goes back to Bernard Williams and Thomas Nagel.¹ The problem turns on our strong intuition to the effect that "people cannot be morally assessed for what is not their fault, or for what is due to factors beyond their control".² But the fact is that we

do morally assess people for things that are not their fault, things that happen due to luck. Of the two truck drivers who fail to have their brakes checked, we blame the one who hits a pedestrian more than we blame the other who does not. Yet the fact that there happened to be a pedestrian in the truck driver's way is not something that is under the truck driver's control. Similarly, we punish murder more severely than attempted murder even where the fact that a person does not succeed in carrying out her intentions is due to accidental factors. This is the puzzle of moral luck: the conflict between the intuition that moral assessment should not depend on luck, and the fact that what appears to be sheer luck often does affect how persons are morally assessed. As Nagel puts it:

How is it possible to be more or less culpable depending on whether a child gets into the path of one's car, or a bird into the path of one's bullet? Perhaps it is true that what is done depends on more than the agent's state of mind or intention. The problem then is, why is it not irrational to base moral assessment on what people do, in this broad sense? It amounts to holding them responsible for the contributions of fate as well as for their own.³

Thus, according to Nagel, our intuition that luck should not make a moral difference is absolutely central to the notion of morality. At the same time, however, Nagel acknowledges the fact that luck sometimes makes a difference to the way a person is assessed morally. If Nagel is right about both of these claims, then the notion of morality is truly paradoxical.

Solutions to the Problem of Moral Luck

There have been two kinds of attempts to resolve the paradox. Some have argued that cases of moral luck are paradoxical only because our notion of morality is flawed.⁴ Browne, for example, has argued that all of our acts have a kind of contingency.⁵ So if moral agents are to be blamed at all for their acts, then they must be blamed for things that happen due to luck. Blame is correctly attributed when an agent is to be held morally responsible for a consequence of her act. The lucky driver cannot be held morally responsible for the death of a child, but the unlucky driver can and should because she killed a child, and this requires a response. We might still believe, Browne admits, that it is *unfair* to blame the unlucky driver more than the lucky driver just because there happened to be a pedestrian in her way. But, according to Browne, we are inclined to hold this belief only because we believe it is wrong to punish people for things that happen due to luck. If we deny that punishment is deserved whenever blame is correctly attributed, the fact that blameworthiness depends on luck is not paradoxical. It is a mere fact that how much a person is to *blame* turns on accidental factors.

The problem with the Browne solution is that it explains away the problem of moral luck. The problem of moral luck, I take it, is not *merely* that no act is complete free of contingency, but also that an agent is blamed *more* for the *same* reckless or evil act than another agent, even though the only difference between the two acts is that one *happens to* result in harm. Even though no act is completely without contingency, we can still compare different degrees of contingency, and the problem of moral luck is that different degrees of contingency can make a moral difference. Suppose two drunk drivers both hit a child in two almost exactly similar situations. Suppose the only difference between the two situations is this. In the one situation, the child dies in the hospital. In the second situation, the child is in intensive care for two days but survives. That the second child survives is due to contingent factors. Perhaps the second child is

slightly stronger physiologically. If Browne were right, then it would be *mistaken* to blame the fortunate drunk driver to the same degree as the unfortunate drunk driver, even though their actions are only different because of a difference made by chance. But why must we absolutely take the physiological strength of the child into consideration in determining how much blame a drunk driver should receive?

The other sort of response to the problem of moral luck consists in denying that how much a person is to blame ever depends on luck (I shall refer to this as the ‘anti-luck theory of blame’).⁶ There is, on this view, no difference between the moral status of the lucky and unlucky drivers; but in the case of the unlucky driver, we are in a better epistemic position to evaluate her defective moral character. The unlucky event reveals the driver’s character (it reveals that she is the kind of person who does not bother to check her brakes). The character of the two drivers, from the point of view of blame, is however exactly the same. As Zimmerman puts it:

it must be admitted that, from the point of view of describing moral responsibility, it does not matter whether or not the terrible event—the death—comes about as a result of the decision to drink and drive, as long as the decision itself occurs. But, again, this does not imply that the successful driver is not to blame for the death. The death indicates the need for a negative evaluation only indirectly, as it were, while the decision indicates it directly; but an indirect indication is still an indication. While the death that results is not itself the occasion of a fresh negative evaluation, it nevertheless reflects the fact that some negative evaluation is called for.⁷

This response to the problem of moral luck is appealing. The main problem with it, however, is that, if it is true, then most of our ordinary moral assessments will turn out to be mistaken. We do after all blame the unfortunate driver more than the fortunate driver. If Zimmerman's solution is correct, then our common intuitions about who is to blame and to which degree are nearly always mistaken, and common-sense ethics must be rejected as flawed. But are we willing to take such a drastic step to solve the problem of moral luck? As Brandt (1958, p. 6) has pointed out, there are only two reasons for analyzing moral concepts: either we are advocating a definition of what people actually mean by their ethical terms, or the actual meanings are somewhat vague and the definition in question is a good suggestion for a modified use of the terms that is rather close to the vague actual meanings. Suggesting as Zimmerman does, that we are simply and completely in error about the way we use ordinary terms—like 'blame' or 'punish'—appears to suggest a complete revision of our moral and legal discourse.

On the anti-luck theory of blame, we are mistaken even if we blame ourselves more when an act of negligence has harmful consequences. Wolf who defends an anti-luck theory of blame provides the following solution to this problem of self-blame.⁸ She argues that people who are equally negligent are equally blameworthy. The unfortunate and the fortunate truck drivers *deserve* the same amount of blame. Yet the drivers should not blame themselves to the same degree. That is, we would expect the unfortunate truck driver to blame herself more than the fortunate truck driver. But this does not mean, Wolf argues, that one of the truck drivers is mistaken, or that we are mistaken when we have different expectations as to how much they should blame themselves. Rather, Wolf argues, it is a virtue when a person can realize that the way things turned out is causally connected with her acts. Realizing the existence of such a causal connection should be manifested in a special kind of regret: regretting the way things

turned out (given one's act) rather than regretting one's act as such. According to Wolf, the two truck drivers *deserve* to be blamed equally by us, but the unfortunate truck driver ought to blame herself more than the fortunate truck driver.

The main problem with this approach is that it does not provide a solution to the problem of why other people, too, tend to blame the unfortunate truck driver more than the fortunate driver. Wolf's view is thus faced with the same problems as Zimmerman's.

Moral Skepticism

One might respond to the criticism of the anti-luck theory of blame by saying that often enough ordinary people are mistaken when they use certain concepts. Many thinkers hold that 'S sees O' implies 'O exists'. But in ordinary life, we are often quite happy to say 'what S saw didn't exist'. Similarly, knowledge is consistent with belief, but we often say, 'S didn't believe it, he knew it'. In each case, we must say that the ordinary judgment is mistaken. In each case, we can do this by saying that our judgments ignore subtle differences that we can recognize under reflection. We wouldn't have to reject ordinary theories of perception to accommodate our commonsense claims such as 'MacBeth saw a dagger that didn't exist'.

This objection is well taken. There is, however, a further problem with the anti-luck theory of blame. Consider again Zimmerman's unfortunate drunk driver who kills a child. Surely, a drunk driver is very much to blame for his act. But consider now the unfortunate drunk driver's counterpart who has exactly the same defective character traits as the unfortunate drunk driver and who gets drunk and gets into his car to drive home but whose car fails to start. On the anti-luck theory of blame, the fortunate drunk 'driver' who did not drive drunk because the car wouldn't start is as much to blame as the unfortunate driver who killed a child, even though he

did not drive drunk. For had the car started and had a child run out in front of the car, he would have killed a child. Consider now the unfortunate drunk driver's counterpart who has exactly the same defective character traits as the unfortunate drunk driver but whose car breaks down before he gets to the party and who therefore does not even get drunk that night. He, too, is as much to blame as the driver who kills the child, even though he was not even drunk. For had he gotten drunk and had his car started and had a child run out in front of the car, he would have killed a child. But then consider the unfortunate drunk driver's counterpart who is born with exactly the same mental capacity as the unfortunate drunk driver but whose morally defective parents die when he is a baby and who therefore never develops bad character traits. Shouldn't he also be as much to blame as the driver who kills a child? After all, it isn't the unfortunate drunk driver's fault that he grew up with morally defective parents.

If the answer is in the negative, then it is allowed that agents are blamed for things that happen due to luck. The unfortunate drunk driver's counterpart who has exactly the same defective character traits as the unfortunate drunk driver but whose car breaks down before he gets to the party should be blamed more than the counterpart who never develops bad character traits, even though the only difference between them is the difference made by chance. But once we admit that the amount of blame an agent deserves depends on luck, there is no longer a principled reason for denying that the drunk driver who kills a child deserves more blame than the drunk driver who does not kill a child. On the other hand, if the answer to the above question is in the affirmative, the notion of blame is truly counterintuitive. We should then sometimes blame people who do not have defective character traits because they have counterparts with defective character traits. A more plausible stand is this: that luck never plays a role in how much an agent is to be blamed. But then the anti-luck theorist is committed to a view according to

which no agent is to blame for anything; for no act is entirely free of contingency. Nagel compares the skeptical consequences of an anti-luck theory of blame to epistemological skepticism. As he says:

It resembles the situation in another area of philosophy, the theory of knowledge. There too conditions which seem perfectly natural, and which grow out of the ordinary procedures for challenging and defending claims to knowledge, threaten to undermine all such claims if consistently applied. Most skeptical arguments have this quality: they do not depend on the imposition of arbitrarily stringent standards of knowledge, arrived at by misunderstanding, but appear to grow inevitably from the consistent application of ordinary standards. ... Our beliefs are always, ultimately, due to factors outside our control and the impossibility of encompassing those factors without being at the mercy of others leads us to doubt whether we know anything.⁹

What Nagel is getting at is this. We have an intuition that beliefs, to be knowledge, must be both non-accidentally acquired and so sufficiently justified that they are unlikely to be false. But the fact is that none (or very few) of our true beliefs satisfies these requirements. Suppose water is pouring down outside my window in a rain-like fashion. I acquire the belief that it is raining. My belief is true: it *is* raining. But my belief is far from being certain. It is after all possible, given the evidence at my disposal, that someone upstairs is playing a trick on me with a rain-machine. Since I didn't check whether or not somebody is indeed playing such a trick, I am to some extent

merely *lucky* that no one was in fact doing so. I am *lucky* that my belief turns out to be knowledge.

The common-sense notion of knowledge is however such that we do hold that I know that it is raining under the given circumstances, even in spite of my being lucky. On the common-sense notion of knowledge, I need not be maximally justified in believing that it is raining in order to have knowledge of this fact. When, however, we reflect on the notion of knowledge, then we very quickly come to the conclusion that a given belief ought to count as knowledge only if the evidence we have is enough to guarantee that that belief is also true. There is a conflict here between the intuition that knowledge should be non-accidental and certain and the fact that our everyday knowledge claims rarely meet these two requirements.

Suppose that we were to solve this conflict in a way similar to the Zimmerman solution to the problem of moral luck.¹⁰ This would mean affirming, not that our intuition that knowledge must be completely certain and non-accidental is mistaken, but rather that we are wrong about our everyday knowledge claims. I am wrong to think that I know that it is raining, even when my belief is true and I can see that water is pouring down.

To reject all of our everyday knowledge claims in this way is to embrace skepticism. And similarly: those who argue that we are wrong when we blame the murderer more than the person who merely attempted murder are, I shall propose, embracing some form of moral skepticism. In moral skepticism, as in epistemological skepticism, only the completely non-accidental counts. In moral skepticism, only the completely non-accidental features of an agent's mental and physical acts fall into the domain of what can be morally assessed. In epistemological skepticism, only the completely non-accidental belief can count as knowledge. Epistemological skepticism has the consequence that we cannot really know anything, because we are never

completely justified in our beliefs. Moral skepticism has the consequence that we cannot really be to blame for anything, for how can we guarantee that any of our acts, whether mental or physical, are completely free from being influenced by accidental factors?

The Problem of Generality

If we are ever to blame for our acts, the anti-luck theory of blame must be false. So perhaps we should reject it and embrace instead a form of act consequentialism according to which how much a person is to blame depends on how wrong her intended act is.¹¹ This approach is appealing. Unfortunately it is not without its problems. In fact, act consequentialism, I shall argue, gives rise to the problem of which level of generality the act for which an agent is blamed should be described. The ethical problem of generality is analogous to the epistemological problem of which level of generality a reliable method for acquiring knowledge should be described at.

The problem arises once we consider what ‘wrong act’ means. Suppose the alleged wrong act of *killing a child* is an example of a wrong act on this view. If this were so, then we could argue that the two truck drivers carried out different wrong acts. So the drivers ought to be blamed differently (other things being equal). But this would beg the question against the anti-luck theorist. Whatever ‘wrong act’ means on this view, it does not refer to a willful act and its (actual) harmful consequence. ‘Wrong act’, I take it, must refer to a willful act that has a *potential* harmful consequence. Loosely speaking, it is a willful act to drive without having your brakes checked just in case you have not taken any measures to prevent the act. It is furthermore wrong to drive without having your brakes checked, because driving under those circumstances is likely to inflict harm. A truck driver is thus to blame if she is driving without having her

brakes checked, for she is carrying out a wrong act thereby. On the present consequentialist theory of blame, it does not matter whether or not she in fact inflicts harm but only whether or not it is likely that her failure to check her brakes will do so.

The present consequentialist theory also has built into it a ranking of blame. The more likely it is that an act will inflict harm (and the more severe the harm)¹², the more the agent is to blame. For each kind of harm, the amount of blame that is assigned to an act is a function of how likely the act is to inflict that harm. Suppose John scratches his friend Peter's arm in anger. Peter suffers from a skin disease, he is inflamed by the scratching, and he goes to Florida to recuperate. In Florida, he is run over by a train and is killed. Surely, his death is a consequence of John's act. Indeed whenever one person scratches another, it is possible that death ensues. But there is no doubt that the amount of blame John should receive for his act of scratching Peter's arm should be very small. For it is not very likely that a person dies as the result of a scratch on his arm.

Since the amount of blame a person should receive depends on how *likely* it is that her act will inflict harm, this amount cannot be measured by the act itself. It is most certainly act types that are likely to inflict harm, not act tokens. Yet people are to blame for act tokens, not types. So, in order to decide whether a person is to blame for an act token, we must identify the act type to which it belongs. But each act token belongs to more than one type, and different act types are associated with different degrees of likelihood. The different act types will vary in respect of likelihood to inflict harm depending on how broad or narrow they are. An act type X is more narrow than an act type Y just in case necessarily every instance of X is an instance of Y, but not *vice versa*. For each act type of which a given act token is an instance, there is a matching description of that token. The more detailed the description of the token, the more narrow the type under which it is classified, and the more general the description, the broader the type. The

ethical problem of generality is that there seems to be no principle that determines the right level of generality: the more likely an act is to inflict harm, the more blame a person should receive for that act.¹³ But how likely an act is to inflict harm depends on the level of generality specified. Since no one level is better than any other, it would seem that there is no non-*ad hoc* way to specify the correct level of generality at which to describe an act when appraising how much blame its agents should receive. Scratching a person's arm is unlikely to result in the death of that person. This might not be so in the case of scratching a person who has a severe skin condition and who is likely to need to go to a place with heavy traffic in order to recuperate. At one level of generality John's act is not likely to result in severe harm. At another level of generality, it is. How much blame John should receive for his act depends on how likely it is that his act inflicts severe harm, but the latter depends on the level of generality at which the harmful act is described.

Epistemological Contextualism

I will suggest that if we solve the problem of generality, then we have also solved the problem of moral luck. The problem of generality arises, I will argue, because we ignore that the extension of the term 'blame' is determined by context. The present consequentialist theory of blame is in this respect similar to process reliabilist theories of knowledge.¹⁴

The more familiar reason that contextualism is advocated, however, is that it appears to solve the problem of skepticism. As Heller has argued, the fundamental intuition underlying reliabilist theories of knowledge is this:¹⁵ that in order for a subject to have knowledge, there must be a necessary connection between what makes her belief true and why she acquired it. In

other words, the subject's belief must be so well connected with the facts—and so her means of acquiring the belief so reliable—that there is no possibility of it being false.

The way to maintain the 'anti-luck' variation of reliabilism without embracing epistemological skepticism is to combine it with contextualism:¹⁶ the subject's belief must be *so* well connected with the facts that there is *no possibility* of it being false, but it is the interests or inclinations of the conversationalists that establish what counts as a false belief-possibility in any given case. When Mark wrapped up the leftovers last night, he carelessly put the roll of the plastic wrap in the refrigerator and the leftovers in the drawer. Mark's wife stumbled on the wrap in the refrigerator, worked out what happened, and placed the wrap in the drawer and the leftovers in the refrigerator. Today both of them believe that the leftovers are in the refrigerator, and neither one of them is concerned with whether or not they are hoodwinked by an evil genius. They *are* interested, however, in preventing the food from rotting. Mark's belief might just as easily have been false. His wife's belief is better connected to its truthmaker. It is sufficiently well connected with the truth in this context to comprise knowledge. Mark's wife knows that the leftovers are in the refrigerator; Mark does not.¹⁷

If a skeptic came along, one seriously entertaining the possibility that there are evil geniuses afoot, then her conversational interests would conflict with those of Mark and his wife. She, accordingly, would be in a different context, a context in which both would fail to know that the leftovers are in the refrigerator. She might even convince Mark and his wife to switch their conversational interests in such a way as to fit her own. Both would then agree, in their new context, that even Mark's wife fails to know that the leftovers are in the refrigerator.

It is important to notice that on a contextualism of this sort, the subject remains the same no matter which context is chosen. Mark's wife has the same true belief and the same evidence

in both her ordinary context and in the skeptic's context. Whether or not we can correctly apply the term 'knowledge' to her true, non-maximally justified belief that the leftovers are in the refrigerator varies, not with any of the epistemic properties she possesses, but only with the conversational interests or inclinations of the person using the term.

Where do the interests and inclinations derive from? As Lewis has suggested in his "Scorekeeping in a Language Game", the interests and inclinations come from presuppositions, habits, and conventions, in conjunction with facts about what the world is really like. While habits and conventions usually change only over a long period of time, presuppositions can be created or destroyed in the course of a conversation. Usually, when Mark and his wife enter a conversation, they presuppose that there is a world, that the world exists more or less as they experience it, that their perceptual systems function optimally, that the light in the room is normal, and so forth. But if a skeptic enters the conversation and brings forth the possibility that they are deceived by an evil genius, then these presuppositions are destroyed. Of course, the skeptic might not succeed in changing the context of Mark and his wife, but they are destroyed at least in her context. She cannot talk with Mark and his wife about skeptical possibilities without bringing a new set of presuppositions into existence. As Lewis says,

it's not as easy as you might think to say something that will be unacceptable for lack of required presuppositions. Say something that requires a missing presupposition, and straightway that presupposition springs into existence, making what you said acceptable after all.¹⁸

We cannot, for example, assume the world we experience exists independently of our beliefs and conceptions and at the same time sincerely entertain the possibility that we are brains in a vat. For mentioning certain skeptical possibilities requires a new presupposition, namely that the world we experience might be a world consisting only of delusions in our brains. So the skeptic cannot talk about brains in a vat without destroying (at least in her context) our most cherished platitudes about the world.

The upshot of this is that our use of the term ‘knowledge’ is such that the very same set of epistemic properties possessed by a person is correctly referred to by our use of the term in one context, but not in another. And what determines the correctness of our use of the term ‘knowledge’ is our inclinations and interests at the time at which we are speaking.

The Context Sensitivity of ‘Blame’

I suggest that we solve the problem of moral luck in much the same way in which the contextualist customarily proposes to solve the problem of skepticism. We can do this by combining the present consequentialist theory of blame with contextualism. So where an anti-luck theorist like Zimmerman would have it that our knowledge of the driver’s blameworthiness varies with the subject’s circumstances (the consequences of her acts), I shall argue that it is the unfortunate driver’s blameworthiness—not just our knowledge of it—that varies with context, and in this case with the context of the evaluator.

Recall that what we are seeking is a means to establish a plausible answer to the question of how *likely* an act is to inflict harm (here again I presuppose that the agent has no good excuses for carrying out the act). In order to determine whether or not a given instance of an act type is

likely to have a certain consequence, the evaluator must consider all relevant possible situations in which acts of this same type are performed.

What makes a possible situation relevant? A possibility is relevant, first of all, if it is a realistic possibility. As Heller has pointed out, “it seems clear that ‘realistic possibility’ is closely connected to such modal terms as ‘likely’ and ‘probable’ and such non-modal terms as ‘rare’ and ‘typical’ ”.¹⁹ One possible event is more realistic than another if it is more akin than the other to the kinds of events that are commonly manifested in the actual world. It is not a realistic possibility that a person dies merely because his friend scratches his arm in anger. This sort of event sequence is extremely rare in the actual world. It rarely happens that someone whose arm is scratched needs to recuperate. It is even more rare that someone who recuperates because his arm is scratched is run over by a train. Because this sort of event sequence is so rarely instantiated in the actual world, there is no realistic way for someone to die merely because his arm is scratched.

Similarly, it is not a realistic possibility that a truck driver who has not had her brakes checked is driving only in areas with no people or cars, or that a truck driver who has not had her brakes checked and runs over what appears to be a pedestrian was in fact running over a dummy, or that good-spirited aliens have installed super-brakes in all the trucks on earth. The combinations of such events are so rare in the actual world that their conjoint occurrences are not realistic possibilities. But it is a fairly realistic possibility that a driver who has not had her brakes checked recently kills a pedestrian who comes in her way. For it happens at least once in a while that brakes, which have not been checked, fail in a situation where the driver needs to stop the car immediately.

Thus all relevant possibilities must be similar enough in *the relevant respects* to the actual world. Even if we limit the relevant possibilities to the realistic possibilities, however, it is still not established what those relevant respects are in which a possibility must be similar to the actual world. As with knowledge, which of the realistic possibilities count as (relevant) possibilities is determined by the interests or inclinations of the conversationalists. That is, it is the conversationalists' context (their presuppositions, shared conventions, habits, and so on) that determines how much significance is to be given to the different relevant respects in which one situation might resemble another. How likely it is that a person who is driving without having her brakes checked would run down a pedestrian depends on how much weight is to be given to the driver's behavior, the effectiveness of her brakes, the environment in which she is driving, and so forth. If the relevant possibilities include only those in which the environment is a school district, then it is much more likely that she kills someone than if the relevant possibilities include only possibilities in which the environment is a silent road on the countryside. If the relevant possibilities are those in which the driver's brakes function optimally despite the lack of a check-up, then it is much less likely that she will kill someone than if the relevant possibilities are those in which the driver's brakes are very ineffective.

Thus, how likely it is that a negligent truck driver kills a pedestrian depends on which possibilities, in which a negligent truck driver is driving, we take to be relevant, given our current context. Different contexts will give rise to different categorizations of possible situations. If we attempt to evaluate whether a given negligent truck driver who killed a pedestrian is to blame, and we focus on her driving in a school district, then the relevant possibilities would tend to include only those in which the negligent truck drivers are driving in school districts and to exclude those in which they are driving on empty roads on the country

side. In this context it is more likely that a negligent truck driver kills a pedestrian than it is in the context focused on negligent drivers on the countryside. Since it will be more likely in this context, more blame ought to be assigned to the agent than if she had killed a pedestrian while driving on the countryside.

Similarly, if the relevant possibilities for determining how likely it is that driving with ineffective brakes leads to the death of a pedestrian include many more possibilities in which a pedestrian runs out in front of the vehicle than possibilities in which no pedestrian runs out in front of the vehicle, then the judgment ought to be: very likely. If the relevant possibilities include less possibilities in which a pedestrian runs out in front of the vehicle and many more possibilities in which no pedestrian runs out in front of the vehicle, then the judgment ought to be: not very likely. In the former context, the driver ought to receive more blame than in the latter.

Moral Luck and Context Switch

Consider again the unfortunate truck driver. Call her 'Unlucky'. When Unlucky's defense lawyer who has studied the statistics looks carefully at the situation in question, he evaluates 'likelihood' with respect to a wide range of possible situations which includes situations in which a pedestrian does get in the way of the truck and situations in which a pedestrian does not get in the way, situations in which the brakes are ineffective and situations in which they are not, and so on. Suppose that when all of these possible situations are taken into account, the likelihood of an accident is low. The lawyer can then truly judge that Unlucky is only slightly blameworthy.

But consider now the parent of the dead child. The parent's focus will likely be switched to a much narrower range of possibilities, which involve relatively more possibilities with a

pedestrian in the way of a truck with ineffective brakes. The likelihood, now, that a truck with ineffective brakes will run down a pedestrian is higher than in the lawyer's context of evaluation. In the context of the parent the judgment 'Unlucky is only slightly blameworthy' is false.

Of course, it may happen that the possibilities that no pedestrian is in the driver's way or that the brakes work despite not having been checked recently is brought to the parent's attention. If the parent says: 'The unfortunate driver should be blamed considerably. She did after all kill my child', we can imagine the defense lawyer saying: 'The fact that your child ran out on the street and that the brakes of the truck did not function optimally is a tragedy. But think of all the drivers who miss a service appointment but whose brakes are effective nonetheless, or who for some other reason are never in Unlucky's situation'. The lawyer thereby brings into consideration a different set of relevant possibilities. She has, in Lewis's terminology, "changed the score on us" (1979, p. 245). When the lawyer thinks that the unfortunate driver deserves a relatively small amount of blame, what he says is acceptable only under higher ethical standards for blame.²⁰ The parent may now switch to the lawyer's context where a different set of alternative situations is used to calculate the likelihood. In the new context of evaluation, the parent ought to truly judge that Unlucky is only slightly blameworthy.

The present contextualist theory of blame solves the problem of moral luck in the following way. The problem of moral luck is, as mentioned, that ordinary people seem to blame the unfortunate truck driver more than the fortunate. I believe that they are right to do so. Ordinary people only very rarely would be in a situation where they would have to blame both an unlucky and a lucky driver who have not had their brakes checked recently. They are rarely evaluating the two in one and the same context.²¹

That is, in ordinary situations we usually judge either an unfortunate driver or a fortunate driver. We might be aware that our good friend who is a truck driver sometimes is a bit sloppy in remembering to have her brakes checked. We are somehow aware that her slightly ineffective brakes could result in an accident. But we do not blame her very much. Since the fortunate driver is fortunate, we will tend to evaluate the situation from within a context in which more of the possibilities in which a pedestrian comes in the driver's way or in which the brakes do not work effectively are ignored. And the likelihood of an accident calculated on the basis of these possibilities is low. So in this context the judgment 'Lucky is only slightly blameworthy' is true.

The case of the unfortunate truck driver is now very different, since a dead pedestrian is involved. Our focus will likely be switched to a different set of possibilities, which involves more possibilities with a pedestrian in the way of a truck with ineffective brakes. The likelihood, now, that the truck will run down the pedestrian is higher. So in this context the judgment 'Unlucky is only slightly blameworthy' is false.

Of course, if we evaluate the moral situation from within one and the same context, then the present contextualist theory commits us to the view that the drivers should receive an equal amount of blame. For if the willed acts of the two drivers are described at the same level of generality, then there is no difference between them in respect of how likely they are to lead to harmful consequences. Both drivers failed to check the brakes under the same set of circumstances. So when 'Lucky is only slightly blameworthy' and 'Unlucky is only slightly blameworthy' are evaluated in the same context, they have the same truth-value.

Though I have argued that there is a parallel between the context-sensitivity of 'blame' and the context-sensitivity of 'knowledge', I do not wish to maintain that the parallel is perfect. In the case of epistemological contextualism the mere consideration of a skeptical hypothesis

forces it to become relevant. If a skeptic makes us consider a skeptical hypothesis, for example by mentioning it, we can no longer truly attribute the knowledge we truly attributed before we considered the hypothesis. So here it seems to be relatively painless to get an evaluator to extend the range of possibilities she takes to be relevant. This is not so, I believe, in the case of moral contextualism. The mere mentioning of the possibility of a negligent driver who avoids accidents, I am sure, is not enough to make the grieving parent switch to the context of the defense lawyer. The parent will most certainly evaluate the fortunate driver under his higher ethical standards for blame. The parent and the lawyer will then disagree on how much blame negligent drivers should receive, but the parent can now truly attribute the same amount of blame to both unfortunate and fortunate drivers.

Though there is a difference between the epistemological and moral context switches, however, moral contextualism is no less contextualist than epistemological contextualism. After all, the heart of contextualism is exactly this: that one and the same epistemic/moral judgment can correctly receive different truth-values in different contexts of evaluations.

Why Aren't Ordinary People Mistaken?

An objection here arises. The defense lawyer and other people who render more possibilities relevant than the parent of the dead child when evaluating the claim 'Unlucky is only a little blameworthy' will evaluate the two drivers equally and not harshly. But why is this contextualism rather than merely the difference between a true ethical judgment and a false one? Why not just say that the lawyer has knowledge that the parent lacks? After all, the lawyer might know more than the parent about the likelihood of accidents given certain acts.

There are two issues here. One is whether the parent who would tend to blame the fortunate and the unfortunate drivers differently is mistaken. The other is whether the parent is mistaken when he blames the unfortunate driver more harshly than, say, the defense lawyer. On the present consequentialist theory of blame, however, the two issues are related. It is true that the lawyer might have better knowledge of statistical correlations. He might indeed consult the statistics when he seeks to determine how probable lethal accidents are when you drive without having your brakes checked. But those in charge of the statistics have already chosen a particular level of generality at which an act is to be described, and so they have already decided which possibilities are to be relevant in determining how likely the act is to have harmful consequences. If there is no principled way to select the proper level of generality at which to describe an act when evaluating how much blame its agent should receive, then an agent may correctly receive different amounts of blame for the same act when the act is described at different levels of generality. Similarly, even though the unfortunate and the fortunate drivers carried out similar act tokens, they too may correctly receive different amounts of blame for their act when the acts are described at different levels of generality.

This is not to say that the same person correctly ought to receive a long imprisonment in one context of evaluation but only a small fine in another. Even context-relative terms are not so dependent on context that just anything goes. The reason that ‘blame’ is context-sensitive is that it is a vague term. It is a vague term just like ‘rich’, ‘tall’, and ‘red’. There are clear-cut cases of people who are rich and poor. Bill Gates is rich. My indigent sister is poor. There are no two contexts such that the latter is rich in one context and the former poor in the other, even though ‘rich’ is vague. But there are also borderline cases. Perhaps an ordinary assistant professor is a borderline case on the scale of rich and poor. But if a given person counts as rich in one context,

then in that same context one who owns more than her is rich too. Such constraints on what is required if a vague term is to be applied correctly are what Fine has called ‘penumbral connections’.²² Fine distinguishes between internal and external constraints. Internal constraints restrain the boundaries of the range of applications of single predicates like ‘rich’ or ‘red’. External constraints confine the shared boundaries of the range of applications of several predicates: even though people in different contexts may disagree on whether something is ‘red’ or ‘pink’, this should not, properly, lead them to consider the same thing as both in one and the same context.

In the case of blame, there are internal constraints on which evaluations are acceptable. For example, if a truck driver is to blame in one context for not having checked her brakes, then every truck driver who has not checked her brakes is to blame *in that same context*. The fortunate truck driver should never, in one and the same context, correctly receive more blame than the unfortunate truck driver; a driver who missed last week’s car service appointment should never, in one and the same context, correctly receive more blame than a drunk driver; someone scratching Peter’s arm should never, in one and the same context, correctly receive more blame than someone stabbing him with a knife. And there are external constraints: even though people in different contexts can disagree on whether or not someone is to blame for a certain act, this should not correctly lead them to count, in one and the same context, the person both to blame and to praise for the act.

Besides violating penumbral connections, an evaluator can, I believe, wrongly attribute blame to a subject who is not to blame in any context, or wrongly fail to attribute blame to a subject who is to blame in any context. If ‘blame’ is a vague term, there will be cases in which it is clearly wrong to attribute blame to a person in any context of evaluation, and cases in which it

is clearly correct to attribute blame in any context of evaluation. Here the parallel between moral and epistemological contextualism holds up. Most epistemological contextualists want to maintain that some contexts are too lenient to be proper contexts of evaluation. Suppose an evaluator occupies a context in which the possibility that one's dreams do not come true is ignored. Unless her context of evaluation is held to be too lenient, the attributor may correctly attribute the knowledge that P to a subject who (truly) believes that P because she dreamed that P. But surely that would render epistemological contextualism much too counterintuitive. The epistemological contextualist ought to require, like Heller above, that no very realistic possibility be ignored. Which possibilities are very realistic depends, not on the context of the evaluator, but on the way the world is.

Undoubtedly, in the skeptic's context not even *very* unrealistic possibilities may be properly ignored. Having accepted a consequentialist theory of blame, there is no analogous situation for the moral case. There is no situation in which a person deserves a lot of blame though it is a very unrealistic (perhaps non-physical) possibility that her act inflicts harm. If it is a very realistic possibility that a given act inflicts harm, then an evaluator who commits it ought to be to blame in any context. Conversely, if it is a very unrealistic possibility that an act inflicts harm, then the person should not be blamed in any context of evaluation. Only in the borderline cases should context play a role. Suppose a person blames a truck driver for not having her brakes checked recently, but that it turns out that the truck driver has a truck with new super brakes that are guaranteed not to fail the first five years of driving. Then surely the evaluator is just wrong in blaming the truck driver for not having her brakes checked no matter the context she occupies. She was wrongly comparing the actual situation to the wrong set of possible situations in which truck drivers without the new super brakes fail to have their brakes checked

frequently. Thus, there are clear-cut cases of a lack of relevant similarity between the actual and the alternative situations.

Of course, in all the borderline cases ‘blame’ is subjective in the sense that its extension varies with the contextual interests or inclinations of those who happen to use it on a specific occasion, and so it is a function of their particular (intentional or non-intentional) choice of a context. Even so, there is no cause for alarm; for which act token a person completes is perfectly objective (just as the height of a person measured in inches is entirely objective, even though whether or not ‘tall’ applies to the person will be a function of the context of those who use the term). That is, what is true of the act token the evaluator is talking about is independent of any linguistic facts about ‘blame’. The agent either completes or doesn’t complete the act token, in the face of any desire, indignation, sorrow, anguish, or other mental attitude of the evaluator. Under which act type the evaluator is sorting the act token when she says ‘she is very much to blame for her act’ depends on semantic facts: it depends on her use of the vague term ‘blame’. It is her use of the vague term ‘blame’ that brings in the subjectivity, not the completion of any given act token that the vague term is predicated of.

Compare the context-sensitivity of ‘blame’ to other context-sensitive words, such as ‘here’. Suppose you are in New Jersey. We are talking on the phone and you truly say about your friend: ‘Peter is here’. When I get off the phone, I truly say to a person in the room who is familiar with the content of our conversation: ‘No, Peter is not here. He is in New Jersey’. Surely, the truth-value of the sentence ‘Peter is here’ will vary with who evaluates the sentence. But which location Peter has is completely objective. The referent of ‘here’ is fixed by the speaker’s situation (in this case, her location), not by Peter’s location at the time in question.²³

And similarly for 'blame'. To say that the truth or falsehood of a moral assessment from the point of view of blame is context-sensitive is *not* to say that the act token talked about as being wrong is context-sensitive. There is an objective fact as to which act tokens the unfortunate and fortunate drivers complete. The indeterminacy lies in our descriptions of the act types to which their act tokens belong. This fact can be captured in theory by giving, along the lines of what Heller has suggested, a conclusive catalogue of all of the acceptable act types of which an act token is an instance together with the degree of likelihood, for each of those act types, that they will lead to harm.²⁴ We could then link each type of act with the corresponding amount of blame that an agent achieving such an act should receive. The complete listing of act types and associated blame is not context-sensitive. There is a practical difficulty in specifying each act token's overall likelihood rating, but it is not this difficulty that gives rise to the context-sensitivity of 'blame'. Rather, the context-sensitivity of 'blame' ensues from the lack of a principled way of determining, independently of a context of evaluation, which act type and associated degree of likelihood from the overall listing of act types is to be associated with an act token in a given situation.

Conclusion

The problem of moral luck is one of conflict between an intuition and a fact. We have an intuition that our use of the term 'blame' should not depend on ('accidental') factors beyond the agent's control. But it is a fact that we do blame people less whose acts fail to have certain consequences, even if the fact that those consequences didn't happen is a matter of sheer good luck.

I have argued that a common solution to the problem of moral luck gives rise to moral skepticism—the view that no one is to blame for her act. If we want to avoid moral skepticism we must say that a person is to blame for a willful act just in case it is likely that the act results in severe harm. The more likely it is that the act results in severe harm the more she is to blame.

Once we accept this consequentialist view of blame, however, it seems that we will also have to accept that the truth-value of sentences used to make moral judgments varies with context. This is because we blame people for act tokens, but we cannot say that an act token is more or less likely to produce an effect. What we should say is that act types are more or less likely to inflict harm. Each act token will fall under many different types of act. And each different act type is coupled with a different likelihood of inflicting harm. Which act type the evaluator associates with the given act token will decide the likelihood and so the blame allocated. And which act type this is will vary with what the evaluator attends to, which in turn varies with her expectations, indignation, pain, anger, and so on, in conjunction with facts about what the world is really like. In other words, it will vary with the context of the evaluator and the world in which she lives.

In the case of the fortunate truck driver, we are likely to classify her failure to check her brakes under an act type associated only with a low degree of likelihood of causing harm. But when we are reminded of the fact that this act type can lead to severe accidents, then we will likely switch to a new context in which the act is classified under a type which is associated with a higher degree of likelihood.

The paradox of moral luck is thus resolved by embracing a contextualist theory of blame. The problem of moral luck arises only if we take no notice of the fact that our ordinary use of the term ‘blame’ depends, for its extension, on the context in which the evaluator finds herself.²⁵

Philosophy Department

Southern Illinois University, Edwardsville

NOTES

¹ See their symposium in B. A. O. Williams and T. Nagel, "Moral Luck" (*Proceedings of the Aristotelian Society* 50, 1976, 115-35; 137-51).

² T. Nagel, "Moral Luck" (Daniel Statman, ed., *Moral Luck*, Albany: State University of New York Press, 1993, 57-71), p. 58.

³ T. Nagel, *Mortal Questions* (Cambridge: Cambridge University Press, 1979), p. 31).

⁴ See e.g. B. Browne, "A Solution to the Problem of Moral Luck" (*The Philosophical Quarterly* 42, 1992, 345-356); M. U. Walker, "Moral Luck and the Virtues of Impure Agency", (Daniel Statman, ed., *Moral Luck*, Albany: State University of New York Press, 1993, 235-250); and J. Andre, "Nagel, Williams and Moral Luck" (Daniel Statman, ed., *Moral Luck*, Albany: State University of New York Press, 1993, 123-129).

⁵ Browne, "A Solution to the Problem of Moral Luck", *op.cit.*

⁶ See e.g. M. J. Zimmerman, "Luck and Moral Responsibility" (*Ethics* 97, 1987, 374-386); D. Lewis, "The Punishment that Leaves Something to Chance" (*Philosophy and Public Affairs* 18, 1989, 53-67); R. Swinburne, *Responsibility and Atonement*, (Oxford: Clarendon Press, 1989), pp. 41ff; H. Jensen, "Morality and Luck" (Daniel Statman, ed., *Moral Luck*, Albany: State University of New York Press, 1993, 131-140); Jensen (1993); B. Rosebury, "Moral Responsibility and 'Moral Luck'" (*The Philosophical Review* 104, 1995, 499-524); and S. Wolf "The Moral of Moral Luck" (manuscript).

⁷ M. J. Zimmerman, "Luck and Moral Responsibility", *op.cit.*, pp. 383-4.

⁸ S. Wolf "The Moral of Moral Luck", *op.cit.*

⁹ Nagel, *Mortal Questions*, *op.cit.*, p. 27.

¹⁰ See D. Statman, "Moral and Epistemic Luck" (*Ratio* 4, 1991, 146-156). Statman argues that epistemic and moral luck are somewhat similar: for the fact that one's belief is true is a good thing and having a true belief may be based on good luck. According to Statman, however, there is an important difference between morality and epistemology: 'While one's status and credit as a knower is not affected by [good] luck, one's status as a moral agent is' (pp. 152-3).

¹¹ As Brandt once argued, whether a person is to blame for a wrong act depends on whether she has a good excuse for carrying out the act. R. B. Brandt, "Blameworthiness and Obligation" (A. I. Melden, ed., *Essays in Moral Philosophy*, Seattle: University of Washington Press, 1958, 3-39). Since in typical cases of moral luck, the agent does not have any good excuses for carrying out the wrong act (I suppose the truck drivers have no excuse for failing to check their brakes), I will here ignore the complications introduced through the excuse component. The blame a person receives for an act, of course, also varies with whether the person *intends* to bring about the harm under consideration or not. We might account for intent by including a description of the agent's intent or lack thereof in our description of the act. Finally, I shall not use the word 'omitted act'. Rather, I shall apply the term 'act' to both acts and omitted acts.

¹² I shall ignore the further complication introduced by the severity component of an inflicted harm.

¹³ Heller makes a series of related points about reliable processes, and the problem of generality facing reliabilism. See M. Heller, "The Simple Solution to the Problem of Generality" (*Nous* 29, 1995, 501-515), p. 502.

¹⁴ For a contextualist solution to the problem of generality, see Heller, "The Simple Solution to the Problem of generality", *op.cit.*

¹⁵ See Heller, “The Simple Solution to the Problem of Generality”, *op.cit.* and Heller “The Proper Role for Contextualism in an Anti-Luck Epistemology” (*Philosophical Perspectives* 13, 1999, 115-129).

¹⁶ Epistemological contextualism has been defended by e.g. S. Cohen, “Knowledge, Context, and Social Standards” (*Synthese* 73, 1987, 3-26); K. DeRose, “Contextualism and Knowledge Attributions” (*Philosophy and Phenomenological Research*, LII, 1992, 913-929); D. Lewis, “Elusive Knowledge” (*Australasian Journal of Philosophy* 74, 1996, 549-567); and Heller, “The Proper Role for Contextualism in an Anti-Luck Epistemology”, *op.cit.*—primarily as a means to solve the problem of skepticism. For a defense of moral contextualism, see P. Unger, “Contextual Analysis in Ethics”, (*Philosophy and Phenomenological Research* LV, 1995, 1-26).

¹⁷ This example is due to Heller, “The Proper Role for Contextualism in an Anti-Luck Epistemology”, *op.cit.*, pp. 118-119.

¹⁸ D. Lewis, “Scorekeeping in a Language Game” (*Journal of Philosophical Logic* 8, 1979, 339-59. Cited as reprinted in *Philosophical Papers*, Vol. 1, New York: Oxford University Press, 1983, 233-249), p. 234.

¹⁹ M. Heller, “Relevant Alternatives” (*Philosophical Studies* 55, 1989, 23-40), p. 34.

²⁰ We can talk of the standards for blame rising so that one is held less and less to blame for what one has no control over, or we can talk of the standards for rightness rising, so that it is harder and harder to behave acceptably. If the former, then it is the lawyer who has higher standards, if the latter it is the parent. Thanks here to an anonymous referee.

²¹ The term ‘context’ should be understood in such a way that even within a single discussion, and perhaps even within an utterance of a single sentence, and evaluator might be ‘within’ two contexts. Perhaps any thought about Lucky will naturally draw attention to more possibilities in which no accident occurs and any thought about Unlucky will naturally draw attention to more possibilities in which an accident does occur. So it may be that within a single discussion we try, intellectually, to keep the level of generality the same for both drivers, but psychologically we are unable to do so. This would explain why the problem of moral luck pulls us in two directions at once. Intellectually we want to judge the two drivers the same, we want to discount the effects of luck, but psychologically we are unable to judge them the same. Thanks here to an anonymous referee.

²² See K. Fine, “Vagueness, Truth and Logic” (*Synthese* 30, 1975, 265-300). Fine uses the internal and external constraints, not in relation to contexts, but rather to regulate which precisifications or sharpenings of vague concepts are permissible.

²³ See DeRose, “Contextualism and Knowledge Attributions”, *op.cit.* p. 925. DeRose uses a similar example in his response to the familiar objection to epistemological contextualism that it entails subject relativism.

²⁴ See Heller, “The Simple Solution to the Problem of Generality”, *op.cit.* p. 513. Heller makes this point about knowledge and reliability.

²⁵ I would like to thank Matt McGrath, Barry Smith, and an anonymous referee for comments on earlier drafts of this paper.