

Moore on Doing versus Allowing Harm *

Carolina Sartorio

University of Arizona

sartorio@arizona.edu

I. Introduction

Michael Moore's *Causation and Responsibility*¹ is a comprehensive and fascinating study of the relationship between the law, morality, and metaphysics. One of the most interesting (and, at the same time, controversial) theses Moore defends in this book is the claim that some central legal concepts are grounded in metaphysical concepts. In particular, Moore emphasizes the key role played in the law by what he argues is the natural relation of causation studied by metaphysicians. A main illustration of this thesis by Moore concerns the distinction between actions and omissions and between doing and allowing harm, which is a pervasive topic in the book (see especially chapters 3, 5, 6, 15, and 18). Moore believes that there is an important metaphysical difference between actions and omissions that is causal in nature: actions can be causes but omissions cannot. This causal difference, Moore claims, generates a moral difference between actions and omissions and between, roughly, doing and allowing harm. (I say "roughly" because, although for Moore the action/omission distinction is at the basis of the doing/allowing distinction, the two distinctions don't always coincide. More on this below.) In turn,

* I am grateful to the participants at the symposium on Michael Moore's book held at the Rutgers-Camden Institute for Law and Philosophy in August of 2010, and especially to Michael Moore for stimulating discussion of his views as well as to Kim Ferzan for inviting me to participate in the event. Thanks also to Juan Comesaña and Christopher Hitchcock for detailed comments on an earlier draft. Some of this material is an expansion of my critical review of Moore's book that is forthcoming in *Mind*.

¹ Oxford: Oxford University Press, 2009.

Moore believes that the moral difference between actions and omissions grounds a difference in legal responsibility, since Moore takes legal responsibility to be grounded in moral responsibility. More particularly, Moore claims that the moral and legal difference between, roughly, doing and allowing harm that arises from the causal difference between acts and omissions generates an asymmetry in the legitimacy of “consequentialist justifications” (justifications in terms of the ensuing consequences). Roughly, it is permissible to allow some harm to occur in order to prevent the existence of more harm, but it is not equally permissible to do harm in order to prevent the existence of more harm.

In this paper I will critically examine Moore’s views on causation and moral responsibility as they concern the distinction between actions and omissions and between doing and allowing harm. Although these are highly contested issues, I won’t dispute the claim that there is a causal difference between acts and omissions or the claim that there is a moral difference between actions and omissions (or between, roughly, doing and allowing harm). But I will challenge the following claims by Moore:

(1) The causal difference grounds the moral difference.

(I will read this as a conditional claim: *if* there were a causal difference of the kind suggested by Moore, it would generate the moral difference that he believes to exist between actions and omissions and between, roughly, doing and allowing harm.)

(2) The moral difference that is grounded in the causal difference generates an asymmetry in the legitimacy of consequentialist justifications.

In what follows I will challenge both (1) and (2). (1) will be the focus of the next section, and (2) will be the focus of section III.

II. The causal difference as the source of the moral difference

On Moore's view, omissions are mere absences: they are absences of events of certain types. Mere absences, he believes, cannot be causes; hence, whereas actions can be causes (in particular, they can cause harm), omissions cannot (ch. 18). Moore believes that this causal asymmetry between actions and omissions generates a moral difference between, roughly, doing and allowing harm. As I anticipated, Moore doesn't believe that the action/omission distinction is identical to the doing/allowing distinction; however, he thinks that they are intimately related. For example, whereas in a "pure" allowing scenario the agent's contribution is straightforwardly an omission, in a "double-prevention" scenario (a central example of an "impure" allowing) it is an action, but the action is related to the harm by means of an absence of some kind (the agent prevents the occurrence of something that would have prevented the harm). As a result, the relation between the agent and the harm is still not causal on Moore's view, just like in a pure allowing scenario, and there is a moral difference between causing a harm and being involved in a double-prevention of a harm (ch. 3). For the most part, in this paper I won't be concerned with the ways in which the action/omission distinction and the doing/allowing distinction come apart. In order to bypass this complication, for the time

being I will restrict my focus to pure (omissive) allowings, and I will use “allowings” to refer exclusively to omissive allowings. Later in the paper the distinction between omissive and non-omissive allowings will become relevant, so I will reintroduce it then (see section III).

Moore thinks that allowing harm is more easily justified than doing harm because doing harm is causing harm and allowing harm is not causing harm. Now, Moore doesn't endorse the claim that agents can *only* be responsible for harms that they cause. He is well aware of the fact that sometimes we hold agents responsible for harms when they merely allow them to happen. For example, we think that a parent has the positive duty to feed her child (to the extent that he or she can) and that this duty is breached if he or she fails to feed the child and the child dies from starvation. The parent is clearly to blame for the death of her child in that case although, on Moore's view, the parent doesn't cause the child's death. Moore believes that there are two distinct desert determiners or bases for (non-inchoate) responsibility: causation and counterfactual dependence. Counterfactual dependence is the relation that obtains between X and Y when it is the case that, if X hadn't occurred, Y wouldn't have occurred either. The parent is responsible for the death of her child in virtue of the fact that the death counterfactually depended on the parent's omission (if the parent hadn't omitted to feed the child, the child wouldn't have died).

The problem that I see with claim (1) is, then, this. Moore wants to claim that the causal difference between acts and omissions grounds the moral difference between doing and allowing harm. However, if causation and counterfactual dependence are two different bases for responsibility, the causal difference between actions and omissions

doesn't generate a moral difference between doing and allowing harm unless this other thesis is true:

Moral difference between causation and counterfactual dependence: There is a moral difference between being causally connected to a harm and being connected to it (only) via counterfactual dependence: being causally connected to the harm makes one more blameworthy, other things being equal.

Now, why think that *this* thesis is true? It is unclear to me whether Moore attempts to give an argument for it and, if so, what that argument is. Perhaps Moore would suggest that it is intuitively obvious. But this is doubtful. Causation, we are assuming, is a relation that is present when there exists something like a physical process (of the right kind) linking two events. Counterfactual dependence, on the other hand, is a "difference-making" relation: Y counterfactually depends on X when X makes the difference between Y's occurring and non-occurring. If an event counterfactually depends on what I did or failed to do then this means that the event wouldn't have occurred *but for* what I did or failed to do. Obviously, this is a very significant relation that one can bear to an event. So, why think that being connected to an event by means of a physical process has more moral significance than making the difference between the event's occurring and non-occurring? It's unclear why we should think so.

Note, in particular, that in Moore's view counterfactual dependence is neither necessary nor sufficient for causation: there are scenarios of causation without counterfactual dependence and scenarios of counterfactual dependence without causation.

So, if there were a moral difference between causation and counterfactual dependence, there would have to be a moral difference between these two kinds of scenario. Do we see such a difference? Is being physically connected to a harm (in a certain way) when the harm would still have occurred in the absence of what one did worse than making the difference between the harm's occurring and non-occurring without being physically connected to that harm? Again, it's unclear that we think this. In fact, I think it is even less clear when the question is put in this way, for this way of posing the question helps to bring out the intuitive weight that the difference-making relation carries for us.

In chapter 18 Moore argues against the significance of counterfactual dependence (relative to that of causation) by appeal to considerations having to do with emotions like guilt. He argues that we don't feel any less guilty when we find out that a harm we caused would still have occurred if we hadn't caused it (pp. 434-435). But, again, I think that this is unclear. Imagine finding out that a harm that you caused would have happened anyway if you hadn't caused it, as in a scenario of causation without counterfactual dependence. Imagine, for example, that you cause someone to die, you feel terrible about it, but you then find out that your victim would have died anyway of natural causes a few minutes later, and there was nothing you or anyone else could have done to prevent that. (This is a kind of "preemption" scenario: the human causes preempt the natural causes.) It doesn't seem crazy to suggest that that would provide at least some relief. In contrast, compare with how you would feel if a harm befalls on someone and this is a harm that you know you (and perhaps only you) could easily have prevented. There would probably be little relief in finding out (say, by reading Moore's work) that omissions are not causes and thus you didn't cause the harm, all you did was fail to prevent it.

My conclusion so far is this. If causation were the *only* desert basis, then the moral difference between doing and allowing harm would flow from the causal difference between acts and omissions. For then causing harm could make us blameworthy and not causing harm couldn't make us blameworthy. However, on Moore's view causation is not the only desert basis (this is a good feature of the view, given its other components; otherwise the view would entail that there are no positive duties, which is very implausible). For Moore, counterfactual dependence is an alternative desert basis in cases of non-inchoate responsibility. But, if counterfactual dependence is a desert basis of its own, the causal difference between acts and omissions doesn't generate a moral difference between doing and allowing harm unless there is a moral difference between causation and counterfactual dependence. And it is unclear that there is such a moral difference.

Moore might respond that an argument for the moral difference between causation and counterfactual dependence is not really needed. In chapter 18 he argues that, if there is a moral difference between doing and allowing harm, then, given the supervenience of the moral realm on the natural realm, there must be a natural difference underlying the moral difference, and the most obvious candidate for that natural difference—the one that provides the best explanation of the moral difference—is a causal difference between actions and omissions (pp. 448-9).

But, is this right? Is the causal explanation proposed by Moore the best explanation of the moral difference between doing and allowing harm? Again, it seems to me that without at least some reason to believe in the existence of a moral difference between causation and counterfactual dependence, the causal difference between actions

and omissions is not even *an* explanation of the moral difference, let alone the best explanation. Moore argues that his proposed causal explanation is the best explanation (or is better than other potential explanations) by comparing it with one main alternative explanation: the view according to which there are two concepts or kinds of causation, one in which omissions can participate and one in which they cannot.² Moore argues that his account is preferable in that it captures what is essentially the same idea—the idea that omissions aren’t “robust” enough to enter in the same kinds of relations as actions—in a simpler way—that is, without multiplying concepts or kinds of causation.

I agree with Moore that the explanation in terms of two kinds or concepts of causation doesn’t do a better job of accounting for the moral difference than Moore’s own explanation (and I also agree that we should avoid multiplying concepts as much as possible). But the reason it doesn’t do a better job is, I think, basically the same as the reason why Moore’s account is insufficient: if both kinds of causation can constitute bases for responsibility (and we’d want to say they can), then appealing to the difference between the two kinds of causation won’t explain the moral difference unless there is some reason to believe that there is a moral difference between the two kinds of causation themselves. At bottom, the problem faced by the view is the same: if one wants to say that the moral difference between doing and allowing harm is grounded in a difference in the way actions and omissions can be related to harms, because actions and omissions enter in different relations with harms, then simply pointing out that they are different

² Someone who endorses this type of view about causation is Ned Hall (see his “Two Concepts of Causation,” in J. Collins, N. Hall, and L.A. Paul, eds., *Causation and Counterfactuals*, Cambridge, Mass: M.I.T. Press, 2004, pp. 225-76). Moore attributes the view to Jonathan Schaffer (in “Causes Need Not Be Physically Connected to Their Effects: The Case for Negative Causation,” in C. Hitchcock, ed., *Contemporary Debates in Philosophy of Science*, Oxford: Blackwell, 2004, pp. 197-216). But I think this is a mistake (Schaffer doesn’t argue for different concepts or kinds of causation and he doesn’t advocate that view in the paper cited by Moore).

relations clearly isn't enough: one needs to show that there is a moral difference between those two relations, if one thinks that both relations constitute bases for responsibility.

By the way, an alternative account of the moral difference would be one that states that the moral difference arises, not from a difference in the way in which actions and omissions can be *related* to harms, but, instead, from a difference in the metaphysical nature of actions and omissions themselves. As I have explained, according to Moore himself there is an important metaphysical difference between actions and omissions: actions are events and omissions are mere absences (this difference, Moore thinks, generates the causal difference between them). So another way of trying to account for the moral difference between doing and allowing harm in naturalistic terms is to appeal directly to this metaphysical difference. On this view, it doesn't matter whether or not the relations in which actions and omissions can enter are the same (maybe they are, maybe they are not; even if they are not the same, they are both bases for responsibility). What matters is the difference in the kind of metaphysical entity that actions and omissions are. On this view, it's simply in virtue of the fact that an action is an entity of one kind and an omission is an entity of another kind (or perhaps not an "entity" at all) that our omissions make us less blameworthy than our actions. (Naturally, one could then extend this account to scenarios of impure allowing, where the agent's contribution is an act but that contribution is mediated by an absence of some sort. Here too the moral difference would be accounted for in terms of the metaphysical nature of certain entities, instead of in terms of the existence of relations of different kinds.)

Note that this explanation doesn't multiply kinds or concepts of causation and, one may argue, is simpler than Moore's own proposed explanation. For, on this view, the

difference in the metaphysical nature of actions and omissions accounts for the moral difference in a more straightforward way than in Moore's view (that is to say, not by means of generating a difference in the relation that actions and omissions can bear to harms) and in a way that needn't commit one to more than one basis for responsibility. Of course, I am not suggesting that a view like this is without its problems (after all, one may ask: What reasons are there to think that the relevant metaphysical difference between actions and omissions or between events and absences is morally significant?). My point is only that a causal explanation, or, more generally, an explanation in terms of a difference in the kinds of *relations* in which actions and omissions can enter, is not the only possible explanation of the moral difference, and it is not at all obvious that it is the simplest explanation. An explanation that appealed to the nature of the *relata* themselves (as opposed to the nature of the relevant relation) would appear to be a simpler explanation of the moral difference.

III. The nature of the moral difference

Now let us assume, for the sake of the argument, that Moore is right and the causal difference between acts and omissions would generate a moral difference between doing and allowing harm. Is he right to think that such a moral difference generates an asymmetry in the legitimacy of consequentialist justifications? That is to say, is claim (2) true?

An obvious problem with that claim is that there are some instances of doing harm for which consequentialist justifications seem acceptable. The most salient example of this is a famous example involving a runaway trolley, which was introduced by

Philippa Foot and Judith Thomson³: a runaway trolley is threatening to kill five people trapped on the main track; an agent could flip a switch and divert the trolley onto a side track, where only one person is trapped. Most people consider flipping the switch to be morally permissible. The justification is at least partly consequentialist in nature (fewer people would die overall), but the agent wouldn't be merely allowing the harm to the one person on the side track by redirecting the trolley; rather, he would be doing harm to him (in contrast, he would be merely allowing the five people on the main track to die if he were not to redirect the trolley). In fact, when Thomson originally introduced the problem⁴, this is actually how she presented it: as a problem for the thesis of the moral difference between killing and letting die (although this is not because she thought that the trolley example proves once and for all that the thesis is false but because she thought that it shows that the thesis cannot be the claim that killing is always worse than letting die, when other things are equal).⁵

Moore has a response to this (ch.3, pp. 75-77). His response is: causation is “scalar,” that is to say, there are degrees of causation, and redirecting the trolley is only “weakly” causing the one's death (causing it to a small degree). In particular, the contribution of the act of redirection is significantly small compared to that of the preexisting threat and the “initiator” of that threat. In virtue of this, Moore actually classifies cases of redirection like this as cases of allowing (non-omissive allowing). But this is odd; after all, there doesn't seem to be any intuitive sense in which these scenarios

³ See P. Foot, “The Problem of Abortion and the Doctrine of Double Effect,” *Oxford Review* 5, 1967, pp. 5–15, and J. Thomson, “Killing, Letting Die, and the Trolley Problem,” *The Monist* 59, 1976, pp. 204–17.

⁴ In her *op. cit.*

⁵ Foot originally thought that the killing/letting die thesis is consistent with the trolley case but this is because she took the agent in her example to be the driver of the trolley as opposed to a bystander, as in Thomson's example. Foot thought that the driver would be killing people either way, by redirecting the trolley and by not redirecting it. But, as Thomson points out, even if this is true, it is not similarly true of a bystander: the bystander kills the one if he redirects the trolley and only lets the five die if he doesn't.

can count as allowings. Moore would probably point to the fact that they share a feature with omissive allowings: the preexistence of a threat. But clearly this isn't enough to make them allowings, since the threat wasn't originally threatening the person who ends up being the victim; hence the agent doesn't *allow* the threat to kill that victim. To me these cases are plainly cases of doing harm, where the harm is done in a peculiar way: by redirecting a preexisting threat.

At any rate, what we call these scenarios is not important. What's important is what Moore claims about them. And what he claims about them is that in those cases (where an agent intervenes by making a "minor course correction" of a preexisting threat) the agent is not a substantial cause of the ensuing harm. And consequentialist justifications, Moore suggests, are legitimate not only in cases in which causation is absent (as in omission cases) but also in cases of weak causation like redirections of preexisting threats (p. 77). Again, then, we have an alleged moral difference that according to Moore can be traced back to a causal difference. This time the causal difference is not the difference between causes and non-causes but that between strong causes and weak causes. So, in order to give proper credit to Moore's response, I think that we should simply bypass the debate about what the doing/allowing distinction amounts to and reformulate Moore's statements of the moral difference and of the underlying causal distinction without making reference to that distinction. What results is the claim that consequentialist justifications can legitimize failing to prevent harm or weakly causing harm but they cannot legitimize strongly causing harm. In other words: it is permissible to fail to prevent harm or to be a weak cause of a harm in order to prevent

more harm from being done, but it is not permissible to be a strong cause of a harm in order to prevent more harm from being done.

Does this solve the problem posed by the redirection scenarios? I will argue that it does not. For the sake of the argument, I will grant Moore the claim that causation comes in degrees, and, also, the claim that redirecting the trolley is only weakly causing the one's death. Still, I will argue that two major problems remain with Moore's proposal.

First, even if redirecting the trolley is only weakly causing the one's death, it *is* causing that death to some degree. Not redirecting the trolley, by contrast, is *not* causing anyone to die, to any degree. Why is redirecting the trolley permissible, then? Moore contrasts the contribution of redirectors with that of the threat itself and with that of "initiators" of threats. He claims that the contribution of a redirector is much smaller than the contribution of the existing threat, and it is also much smaller than the contribution of an event that initiates or launches a threat. But what I am claiming is that the relevant contrast is neither of these, but the contrast between redirecting the trolley and not redirecting it. After all, these are the two available options for the agent in the situation, from which he has to choose.

Moore could try to respond in the following way. For moral purposes (e.g., for the purposes of attributing moral responsibility or blame), being a sufficiently small cause of an event is like being no cause at all. That is to say, beyond a certain threshold, there is no moral difference between contributing to an event to some degree and not contributing to it at all. For all the relevant moral purposes, then, redirecting the train is no more of a cause of the outcome than failing to redirect it. However, this strikes me as extremely implausible. Even if we grant the existence of such causal/moral thresholds, it seems

clear that turning a trolley onto someone (someone who wasn't already threatened by it) is not one of them. Turning a trolley onto someone might be a smaller causal contribution to the person's death than starting the trolley in the first place, but it certainly doesn't seem to be a trivial causal contribution, or one that is on a par with not causing the death at all, for the purposes of attributing moral responsibility or blame. Hence, one important problem that I see with Moore's appeal to degrees of causation in the trolley case is that it is unclear how it helps: even if turning the trolley is contributing to the death of the person on the side track to a lesser degree than starting a new threat altogether, this doesn't seem to be enough to explain why turning the trolley is permissible.

A second problem that I see with Moore's take on the trolley case, and, in particular, with his account in terms of degrees of causation, is the following. Again, Moore suggests that the agent in the trolley case only weakly causes the one's death because all he does is make a small course correction to a preexistent threat. No *new* threat is created; an old threat is diverted from the five to the one. Let us assume that this much is right: the agent only weakly causes the one's death because, instead of creating a new threat for the one, he uses an old threat to cause his death. A serious problem that arises for this suggestion is that there are some cases that share this feature where we think that the agent would be acting *impermissibly*. Here is a famous case of that kind⁶: again, the runaway trolley is threatening the five people on the track, but this time the only way to stop it is to shove a man in front of it, someone who is heavy enough to stop the trolley or to make it derail. Our reaction is that this is clearly impermissible. But note that the agent would be using a preexisting threat in this case (one that is already threatening the five); he wouldn't be creating a new threat. Hence, if the use of an old or

⁶ This is a case by Thomson, in her *op. cit.*

a new threat marks the difference between weakly and strongly causing the one's death, as in Moore's view, it follows that the agent would only be a weak cause of the man's death in this case, and thus killing him is permissible. Again, this strikes us as clearly the wrong result.

Finally, one could try to claim that shoving the heavy man in front of the train is a more substantial cause of his death, even if the agent still uses the same train that was threatening the five, because it requires force on the agent's part, and thus the agent's contribution is not a "minor course correction" like redirecting the train in the trolley case. But, of course, the use of force isn't really necessary: it might be that all that is required is a light push, or the push of a button, or wobbling the footbridge handrail on which he's leaning; still, killing the man seems impermissible. (On the flipside, turning the trolley towards the one might require force on the agent's part such as the yanking of a very heavy lever; still, turning the trolley seems permissible.) Again, then, Moore's appeal to degrees of causation doesn't help explain why it's permissible to turn the trolley. As the example of the heavy man shows, the explanation is not the fact that the agent's causal contribution is reduced significantly because he uses a preexisting threat.

I conclude that the trolley case remains a serious puzzle for Moore's view. However, I think that Moore's general project of finding a causal difference that underlies the moral difference in these cases might still succeed. For I think that there is another way in which the causal difference might be drawn. In the next section I offer a very rough and tentative suggestion about how I think that this could be done.

IV. A different causal account of the moral difference in trolley cases

How could one give a causal account of the trolley cases? For the reasons we have discussed, the proposal is not going to be that turning the trolley is not causing the one's death—since it clearly is—or that turning the trolley is only weakly causing the one's death—since, to the extent that this is true, it doesn't help explain the permissibility of turning the trolley. But one could try to say that there is still a causal feature of the trolley case that makes redirecting the trolley permissible. Namely, although redirecting the trolley is causing a quite specific state of affairs: the *particular* person on the side track dies, it is not causing a more general state of affairs: *someone* in the situation dies.

Let me explain. First of all, note that there are good reasons to adopt a permissive view of the causal relata, in particular, there are good reasons to want to include more general as well as more specific events or states of affairs as possible causal relata. To illustrate this point, consider a simpler single-victim trolley scenario. The victim is trapped on the tracks up ahead and, again, there is a main track, a side track, and a switch. Unfortunately, however, the two tracks reconverge after the switch, and the victim happens to be located exactly where the two tracks reconverge. As a result, the victim cannot be saved by flipping the switch: flipping the switch only determines whether the trolley will reach the victim from one side or from the other. In this case we probably want to say that, even if flipping the switch causes the death to happen in a particular way or through a particular route, it doesn't cause the death to happen. Flipping the switch causes the more specific event or state of affairs but it doesn't cause the more general event or state of affairs.⁷ Arguably, multiple-victim trolley cases have a similar structure: flipping the switch determines *who* will die (whether the five or the one), and thus it is a

⁷ Alternatively, one could appeal to explanation (instead of causation) to make these distinctions. I won't consider this possibility here.

cause of the particular death that occurs, but it is not a cause of the fact that *someone* in the situation dies. Again, there is a more specific state of affairs that flipping the switch causes and a more general state of affairs that flipping the switch doesn't cause.

It is important to realize that in both cases (the single-victim scenario and the multiple-victim scenario) the reason why flipping the switch is not a cause of the more general state of affairs is not *just* that that general state of affairs would still have occurred if the switch hadn't been flipped. That is clearly part of the explanation, but there has to be more to it, since (as we have seen) causation sometimes comes apart from counterfactual dependence, as in preemption cases. The thought is that "switches" (events like flipping the switch in trolley scenarios) are different from preemptors in that they don't create or help sustain a threat, but are mere redirections of preexisting threats. Arguably, it is this fact, coupled with the absence of counterfactual dependence, that accounts for the absence of causation of the relevant event or state of affairs in switching cases.⁸

I have suggested that flipping the switch in the multiple-victim trolley case is not a cause of the fact that someone in the situation dies, although it is a cause of the death of the particular person on the side track. How could this be part of an account of the moral permissibility of turning the trolley? The first step would be to argue that this is a difference that demarcates all the relevant permissible cases from all the relevant impermissible cases. The second step would be to argue that it is a difference that carries moral significance (or enough moral significance). I am not going to argue for this second step; I will note, however, that it strikes me as quite intuitively obvious that such a

⁸ For a more detailed discussion of the difference between switches and preemptors, see my "Causes as Difference-Makers," *Philosophical Studies* 123, 1-2, 2005, pp. 71-96.

causal difference would carry at least *some* moral significance. As for the first step, it seems to me that the trickiest part is to show that “old threat scenarios” where acting is intuitively impermissible (scenarios where no new threat is introduced but it is still impermissible to act, like the heavy man scenario discussed above) are not switches. In particular, one would have to argue that throwing the heavy man in front of the train is causing someone in the situation to die (in addition to causing the particular death), even if the same old threat that was threatening to kill the five is used to kill the one in that case.

Again, one could try to argue for this by appeal to a comparison with a single-victim case. Imagine that an axe has been launched towards a victim. But (perhaps being ignorant of this fact) some time before the axe reaches the person, an evil agent grabs the victim and shoves him onto the blade of the axe, and the victim dies. Presumably, although the agent doesn’t introduce a new threat in this case, he still causes the victim’s death in this case. (Contrast this scenario with one where the agent slightly interferes with the axe’s trajectory in a way that doesn’t affect the end result: this would be a switching scenario and thus the agent would not be causing the victim’s death in this case.)

Similarly, one could try to suggest, if the trolley is threatening to kill some people in the situation and someone shoves another person in front of the trolley before that happens, he causes someone in the situation to die without introducing a new threat. Just like it is possible to cause a particular death by using an old threat, it is possible to cause the more general state of affairs, the state of affairs consisting in the fact that someone in the situation dies, by using an old threat. Despite their differences in the number of

potential victims, the axe scenario and the heavy man scenario are similar enough that we would expect this resemblance in causal structure to hold.

To sum up, then: there appears to be a causal difference between redirecting the trolley from the five to the one and shoving the heavy man in front of the trolley, thus saving the five. Namely, although in both cases acting in the relevant way causes the one person's death, redirecting the train doesn't cause the more general state of affairs consisting in the fact that someone in the situation dies, but shoving the man in front of the train does. One is then a more significant causal source of death in the heavy man case than in the redirection case. And this appears to be morally significant.

Again, this is all very rough and tentative, but I think it is a promising place to start if one hopes to give a causal account of the moral difference in question. I conclude that a causal account of the famous trolley scenarios might be possible, and it might be exactly what is needed, but it is not the kind of causal account envisaged by Moore.

V. Conclusions

In this paper I have examined two theses put forth by Moore in connection with the moral distinction between actions and omissions and between, roughly, doing and allowing harm: the thesis that the alleged causal asymmetry between actions and omissions generates a corresponding moral asymmetry between actions and omissions and between, roughly, doing and allowing harm, and the thesis that such a moral difference generates an asymmetry in the legitimacy of consequentialist justifications. I have argued that Moore has not convincingly shown that the causal asymmetry he sees between acts and omissions would generate the relevant moral difference. Second, I have argued that the

nature of the relevant moral difference (and, in particular, the role played by consequentialist justifications in the statement of that moral difference) is more complex than Moore's view allows for. Moore attempts to accommodate some apparent problems with his view on the nature of the moral difference by appeal to other causal distinctions, in particular, the distinction between weak causes and strong causes. I have argued that those attempts fail. However, unlike with the claim that the moral difference between actions and omissions is grounded in a causal difference between them, I have suggested that in this case a causal account might be exactly what is needed to explain the moral difference. So I agree with Moore that some causal distinctions are morally relevant. We just disagree about *which* are the morally relevant ones.