

## *A New Asymmetry Between Actions And Omissions\**

CAROLINA SARTORIO  
University of Wisconsin, Madison

### 1. Introduction

Compare the following two scenarios:

*Shooting:* I freely decide to shoot Victim. I pull the trigger and Victim dies. An evil neuroscientist has been secretly monitoring my brain. Had I wavered in my decision to shoot Victim, he would have sent a signal to my brain that would have made me decide to pull the trigger anyway.

*Sharks:* While walking by the beach, I see a child drowning. I think I could jump into the water and save him but I deliberately refrain from doing so. The child drowns. Unbeknownst to me, the water is infested by sharks. Had I jumped in, the sharks would have attacked me and prevented me from saving the child.<sup>1</sup>

In both cases, I couldn't have prevented the ensuing death. However, intuitively, whereas I am morally responsible ("responsible," for short) for Victim's death in Shooting, I am not responsible for the child's death in Sharks. Shooting is a "Frankfurt-style" case.<sup>2</sup> In a Frankfurt-style case, it has been argued, an agent is responsible for something that he couldn't have prevented. Thus, in Shooting, I couldn't have failed to kill Victim, but I am still responsible for his death. For I freely made the decision to shoot and I freely acted on that decision (the neuroscientist didn't have to intervene). By contrast, in Sharks, I am not responsible for the child's death. I freely decided to stay on the shore, and the sharks didn't have to intervene, but, somehow, the fact that I couldn't have saved the child given that the water

was infested by sharks seems to relieve me of responsibility for the death. I am responsible for not *trying* to save the child, and my behavior also shows that I am a bad *person* (that there is something defective with my *character*), but I am not responsible for the child's *death*.<sup>3</sup>

In light of the contrast between cases like Shooting and Sharks, some philosophers have suggested that there is a moral asymmetry between actions and omissions. For we can rephrase the difference between Shooting and Sharks in terms of actions and omissions, as follows. In Shooting, I am responsible for my action of *killing Victim*, even though, given the presence of the neuroscientist, I couldn't have failed to kill Victim. By contrast, in Sharks, I am not responsible for my omission of *failing to save the child*, for, given the presence of the sharks, I couldn't have saved the child. As a result, it has been suggested that the following thesis is true—henceforth, the *Old Asymmetry* thesis:

OA: Whereas an agent can be responsible for an *action* even if he couldn't have done otherwise, an agent cannot be responsible for an *omission* if he couldn't have done otherwise.<sup>4</sup>

In other words, OA claims that, whereas one can be responsible for *acting* a certain way even though one couldn't have failed to act that way, one cannot be responsible for *failing to act* a certain way if one couldn't have acted that way.

In this paper I discuss the question whether there is a moral asymmetry between actions and omissions. I argue that there *is* a moral asymmetry, but it is *not* OA. I offer a new asymmetry thesis, and I argue that this new asymmetry thesis succeeds in capturing the moral asymmetry between actions and omissions that is illustrated by the contrast between Shooting and Sharks.

The plan for the paper is the following. The first part is about OA. As I said, I will reject OA; however, it is important to reject it for the *right* reasons. As we will see, there are people who have rejected OA for the wrong reasons. Those people are confused, not in their belief that OA is false, but in the reasons they have for holding that belief. Thus, in the next section, I explain what the *wrong* reasons for rejecting OA are, and why they are wrong. The discussion in that section centers around the conditions of transmission of responsibility. In section 3, I look into this issue in more detail. I put forth a principle of transmission of responsibility that has causation as an essential ingredient and that serves as a springboard for the next two sections. In section 4, I explain how this principle supports the *right* reasons for rejecting OA, and, in section 5, I explain how it gives rise to a new moral asymmetry between actions and omissions. The new moral asymmetry rests on a *causal* principle according to which actions and omissions have different causal powers. I argue for this principle in sections 6 and 7. As a result, the picture that emerges from the paper is the following:

there is a moral asymmetry between actions and omissions, and this moral asymmetry is the consequence of a causal asymmetry between actions and omissions and the role causation plays in the transmission of responsibility.

A note of clarification. Throughout the paper, I will be assuming that omissions can cause things just as actions can. Thus, my proposal will *not* be that there is a moral asymmetry between actions and omissions because actions can be causes while omissions cannot.<sup>5</sup> On the assumption that omissions *can* be causes, I will argue that the conditions under which they are causes are different from the conditions under which actions are causes, and that this results in a moral asymmetry between actions and omissions.

## 2. The *wrong* reasons for rejecting OA

Recall OA:

OA: Whereas an agent can be responsible for an *action* even if he couldn't have done otherwise, an agent cannot be responsible for an *omission* if he couldn't have done otherwise.

Some people have rejected OA for the wrong reasons. They have reasoned as follows.<sup>6</sup> As we have seen, there are Frankfurt-style action cases in which an agent is responsible for his action despite the fact that he couldn't have done otherwise. Shooting was a case of that type: in Shooting, I am responsible for killing Victim although, given the presence of the neuroscientist, I couldn't have done otherwise. Similarly, the reasoning goes, there are Frankfurt-style *omission* cases in which an agent is responsible for his omission even though he couldn't have done otherwise. This is an alleged example:

*Frankfurt-style omission case* (FSOC): I see the child drowning, I think I can save him by jumping into the water, but I freely decide not to jump in. This time there are no sharks in the water, but the evil neuroscientist is monitoring my brain. Had I wavered in my decision, he would have made me decide not to jump in.

Those who have rejected OA for the wrong reasons would claim that, in FSOC, I am responsible for my failure to save the child, even though (given the presence of the neuroscientist) I couldn't have saved him. Thus, they would conclude that it is possible for an agent to be responsible for an omission when he couldn't have done otherwise, and so OA is false.

I will dub this objection to OA, the *Frankfurt-style objection to OA*, and I will dub someone who raises this kind of objection to OA, a *Frankfurt-style objector to OA*. In what follows, I argue that the Frankfurt-style objection to OA is flawed.

The first thing to notice is that the Frankfurt-style objector to OA cannot just appeal to intuition to support his claim that I am responsible for my failure to save the child in FSOC. For the question that arises naturally is: how can I be responsible for my failure to save the child in FSOC (as the Frankfurt-style objector to OA claims), if (as everybody seems to agree) I am not responsible in Sharks? In both cases, I couldn't have saved the child. Then, why am I, according to the Frankfurt-style objector, responsible for my failure to save the child in FSOC but not in Sharks?

The Frankfurt-style objector will probably try to say the following.<sup>7</sup> In Sharks, had I decided to jump in, I still wouldn't have saved the child (because the sharks would have stopped me). In FSOC, by contrast, had I decided to jump in, I *would* have saved the child (because there aren't sharks or any other obstacles in the water). True, I couldn't easily have decided to jump in, since the neuroscientist was determined not to let me make that decision. Still, *had* I decided to jump in, I would have saved the child. That is, in the closest possible world where I decide to jump in—a world where the neuroscientist fails and where there are no obstacles in the water—I save the child. In other words, in FSOC, whether the child lived or died *hinged on* what I decided to do. And in all of these cases I am responsible for making the decisions that I made (no one forced me to make them). Presumably, then, I am also responsible for my failure to save the child in FSOC. Or so the Frankfurt-style objection to OA goes.

I take it that the thought is that, given that the child's death depended on my decision, for which I am responsible, I am responsible for the child's death. And thus, given that I am responsible for the child's death, I am responsible for my failure to save the child. This seems like a natural thing for the Frankfurt-style objector to say. So I suggest the following reconstruction of the argument by the Frankfurt-style objector:

- (1) In FSOC, I am responsible for my decision not to jump in.
- (2) The child's death depended on that decision.
- (3) If I am responsible for X, and an outcome Y depends on X, then I am responsible for Y.
- (4) Therefore, I am responsible for the child's death.
- (5) If I am responsible for the child's death, then I am responsible for my failure to save the child.
- (6) Therefore, in FSOC, I am responsible for my failure to save the child.

In what follows, I will not discuss step (5). I think it is intuitively plausible that, in a case where I fail to save the child, if I am responsible for the child's death, then I am also responsible for my failure to save him.<sup>8</sup> So I will assume that all the Frankfurt-style objector has to prove to make his case is that I am responsible for the child's death in FSOC.

A note about (3) is in order. As it stands, (3) is not very plausible, for it makes agents come out responsible for things that they couldn't possibly have foreseen. To make it more plausible, we should probably add more provisos; in particular, we should probably say: "If I am responsible for X, an outcome Y depends on X, and it was foreseeable that Y (or an outcome of Y's type) was likely to follow X, then I am responsible for Y." I will not be concerned with any extra provisos because it is likely that they will be met in FSOC (and in other cases I will focus on) and thus, they will not matter for my purposes here. For instance, in FSOC, it was foreseeable that the child's death would likely follow my decision not to jump in. I will assume that (3) can be suitably revised in this respect; in what follows I will refer to the extra needed provisos as "the required provisos."

I will call the argument above an *argument from dependence*. That argument relies heavily on a certain concept of dependence that allegedly links my *decision not to jump in* to *the child's death* in FSOC, and that purportedly transmits my responsibility from one onto the other. I will argue, however, that there is no concept of dependence that can do that, and thus the argument fails. The discussion that follows will center on the issue of the transmission of responsibility. This topic will play a fundamental role in the formulation of the new asymmetry later, so it is important, not only because of its implications for the Frankfurt-style objection to OA, but also because of its implications for the new asymmetry.

What can be the concept of *dependence* that the Frankfurt-style objector has in mind and that purportedly serves to transmit my responsibility to the ensuing outcome? A natural candidate is *causation*. For it is natural to regard causation as a link that agents have with the world, in virtue of which they can be responsible for what happens in it. In other words, it is natural to regard causation as the means by which the responsibility of agents for their actions and omissions transmits to outcomes in the external world. In the next section I will argue that, not only is this a natural picture of how the responsibility of agents is transmitted to outcomes, but it is also the *right* picture. Now, however, I will argue that the Frankfurt-style objector cannot appeal to the concept of causation in his argument from dependence against OA.

If dependence is causation, then the relevant premises of the objector's argument about FSOC read:

- (2') My decision not to jump in caused the child's death.
- (3') If I am responsible for X, and X caused an outcome Y (and the required provisos are met), then I am responsible for Y.

But (2') is false. Arguably, what caused the child's death is not my *decision not to jump in*, but my *failure to decide* to jump in. For the child didn't die in FSOC in virtue of what I *did* decide to do; he died in virtue of what I *didn't* decide to do.<sup>9</sup>

Let me explain. First of all, it is important to distinguish between my decision not to jump in and my failure to decide to jump in. Arguably, these are not identical because they have different modal properties: they would have obtained under different conditions. In particular, my failure to decide to jump in could have obtained without my decision not to jump in obtaining. For instance, I would have failed to decide to jump in if I hadn't made any decision at all but had kept deliberating what to do until the child drowned.<sup>10</sup> The question arises, then: which one of the two, my decision not to jump in or my failure to decide to jump in, has a better claim to cause the child's death in FSOC? Presumably, the latter. For one, the child would still have died if my failure to decide to jump in had obtained without my decision not to jump in obtaining (if I had remained undecided). Also, the fact that I decided not to jump in seems relevant to the child's death *only to the extent that* it entails that I didn't decide to jump in (if I decided not to jump in, I couldn't have decided to jump in, and thus, I couldn't have saved the child). This suggests that what is causally relevant to the child's death is the fact that I didn't decide to jump in, not the fact that I decided not to jump in.

I have argued that premise (2') of the Frankfurt-style argument against OA fails: the child's death wasn't caused by my decision not to jump in, but by my failure to decide to jump in. Now, in light of this objection, the Frankfurt-style objector might think of substituting "my failure to decide to save the child" for "my decision not to save the child" in the relevant premises of the argument, thus:

- (1') In FSOC, I am responsible for my failure to decide to jump in.
- (2'') My failure to decide to jump in caused the child's death.

As I have argued, (2'') is probably true. But now the problem is that (1') is, at the very least, controversial. What *is* clear about FSOC is that I am responsible for what I *decided* (a mental *action* of mine); it is not equally clear that I am responsible for what I *failed* to decide (a mental *omission* of mine).

Let me explain. Premise (1) of the original Frankfurt-style argument was uncontroversial because my decision not to jump in is an *action* and because, as Frankfurt-style *action* cases show, I can be responsible for actions that I couldn't have avoided when I perform them freely (when no one forces me to perform them). But the revised premise, (1'), is *not* supported by a similar reasoning. True, no one forced me to fail to decide to jump in either. But this isn't enough to show that I am responsible for failing to decide to jump in. For, as we have seen, omission cases behave differently from action cases. In Sharks, for instance, no one forced me to fail to save the child, but I am still not responsible for failing to save the child. If the Frankfurt-style objector to OA wants to insist that, in FSOC, I am

responsible for failing to decide to jump in (not just for deciding not to jump in), then his argument becomes question-begging. For there is no more reason to believe that I am responsible for failing to decide to jump in (his premise) than there is to believe that I am responsible for failing to save the child (his conclusion), or, at least, the Frankfurt-style objector hasn't given us any such reason. After all, failing to decide to jump in and failing to save the child are both omissions, and, given the presence of the neuroscientist in the background, I could have avoided neither of them.

In this section I argued that the Frankfurt-style objection to OA fails, on the assumption that causation is the type of dependence that transmits the responsibility of agents for their actions and omissions to outcomes in the world. In the next section, I will argue that this assumption is very likely true. This will close my discussion of the Frankfurt-style objection to OA, and it will also get us a step closer to the new asymmetry.

### 3. The role of causation in the transmission of responsibility

If not causation, then what other type of dependence could transmit responsibility to outcomes? Maybe *counterfactual dependence* could? In what follows, I argue that counterfactual dependence isn't a suitable candidate. Counterfactual dependence will be my specific target, but, as we will see, my argument is likely to extend to other possible candidates as well. I will conclude that causation is the most suitable candidate for transmitting responsibility to outcomes.

Counterfactual dependence is defined in the following way:

Y counterfactually depends on X just in case, had X not occurred, Y would not have occurred.

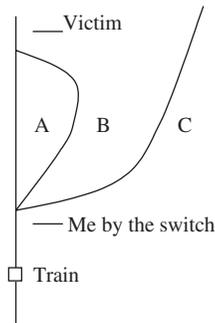
The principle according to which counterfactual dependence transmits responsibility to outcomes is the following:

TR (Counterfactual): If I am responsible for X, and an outcome Y counterfactually depends on X (and the required provisos are met), then I am responsible for Y.

Notice that, if TR (Counterfactual) were true, it could help the Frankfurt-style objector in his argument against OA. For it is possible to fill in the details of FSOC so that the child's death counterfactually depended on my *decision* not to jump in. For instance, we could imagine that I am a very decisive type of person. So, had I not decided not to jump in, I *would* have decided to jump in and thus, I would have saved the child. That is, in the closest possible world where I don't decide not to jump in (a world where the neuroscientist failed, I am a decisive person, and there are no obstacles in the water), I decide to jump in and the child lives. Hence, in FSOC thus

conceived, the child's death counterfactually depends on my decision not to jump in. And I am responsible for this decision. Hence, if TR (Counterfactual) were true, it would follow that I am responsible for the child's death, and then, for my failure to save the child in FSOC. So FSOC would be a counterexample to OA.

But TR (Counterfactual) *isn't* true, as shown by the following case. Imagine that a runaway train is going along the main track (track A) when it approaches a switch. I am by the switch and I see that, if the train continues on track A, it will run over a person, Victim, who is standing further down the track. Although switching to B wouldn't help, switching to C would, as depicted by the following picture:



Imagine that I have an irresistible urge for flipping switches and that I don't care whether Victim lives or dies. I think about where to flip the switch, and I decide to switch it to B. Victim dies. Unbeknownst to me, a neuroscientist has been monitoring my brain and is in control of my bodily movements. Had I decided to switch the train to C, he would have forced my finger into the B-position. So I couldn't have saved Victim.

In this case, given that I freely decided to flip the switch to B, I am responsible for flipping the switch to B (although I couldn't have done otherwise). But I am not likewise responsible for Victim's death. After all, the case has the same structure as Sharks in this respect: I could have decided to switch the train to C, but I couldn't have saved Victim, for the neuroscientist would have stopped me. (Compare: in Sharks, I could have decided to jump in, but I couldn't have saved the child, for the sharks would have stopped me.)<sup>11</sup> But Victim's death counterfactually depends on my flipping the switch to B: had I not flipped it to B (which means that the neuroscientist would have failed), I would have flipped it to C and Victim would have lived. For, in the closest possible world where I don't flip the switch to B (a world where the neuroscientist fails and where I have an urge for flipping switches) I flip it to C and Victim lives. Thus, I am responsible for switching the train to B, I am not responsible for Victim's death, but Victim's death counterfactually depends on my switching the train to B. In other words, TR (Counterfactual) is false.<sup>12</sup>

At the same time, this case serves to illustrate the fact that there can be counterfactual dependence without causation.<sup>13</sup> In this case, Victim's death counterfactually depended on my switching the train to B but wasn't caused by it. If anything, Victim's death was caused by *my failure to switch the train to C*, the safe track. For the fact that I switched the train to the unsafe track is relevant to Victim's death *only to the extent that* it entails the fact that I didn't switch the train to the safe track. This suggests that the cause of Victim's death was my failure to switch the train to C, not my switching it to B instead.<sup>14</sup>

As a result, what emerges from the example is that, when causation is absent, responsibility does not transmit to the ensuing outcome, even if there is counterfactual dependence. Thus, even if I was responsible for switching the train to B, my responsibility does not transmit to Victim's death, because my switching the train to B didn't cause Victim's death. The example also supports the idea that, when there is causation (and the required provisos are met), responsibility does transmit to the outcome. Thus, in the example, *had* I been responsible for failing to switch the train to C, then I would *also* have been responsible for Victim's death, given that my failing to switch to C caused Victim's death. In this case, however, I was not responsible for failing to switch the train to C, because I couldn't have switched it to C (again, the case is like Sharks in this respect: I could have decided to switch the train to C, but I couldn't have acted on that decision because the neuroscientist would have stopped me). Thus, I am not responsible for Victim's death.<sup>15</sup>

I have argued that the most plausible candidate for transmitting responsibility to outcomes is causation. I don't intend this to be a knockdown argument against any other possible candidate, but I think it does make a good *prima facie* case for my view. In a nutshell, the view that I have defended is the following:

TR (Causal): An agent's responsibility for X transmits to an outcome Y iff X causes Y (and the required provisos are met).<sup>16</sup>

TR (Causal) says that causation is necessary for the transmission of responsibility to outcomes (other notions of dependence by themselves won't do), and it says that it is also sufficient, with the required provisos.

In what follows, TR (Causal) will play a double role: it will generate the *right* reasons for rejecting OA, and it will support the new moral asymmetry between actions and omissions. I take these up in turn in the next two sections.

#### 4. The *right* reasons for rejecting OA

As we have seen, FSOC is not a counterexample to OA, or, at least, the Frankfurt-style objector hasn't given us good reason to believe that it is. I submit that, by contrast, the following case *is* a counterexample:

*Planted Sharks*: This time I am responsible for the sharks being in the water: yesterday I negligently released the sharks in the area—I had no good reason to release the sharks, and I had good reason not to, but I still did. Today I see that the child is drowning but I do not attempt a rescue because it would be fruitless.<sup>17</sup>

I will argue that *Planted Sharks* is a counterexample to OA because, although I couldn't have saved the child today, I am responsible for failing to save him today. Thus, it is possible to be responsible for an omission when one couldn't have done otherwise, and so OA is false.

The reason that I couldn't have saved the child in *Planted Sharks* is, again, that the sharks are in the water.<sup>18</sup> And the reason that I am responsible for my failure to save the child is that I am responsible for planting the sharks, which was one of the causes of the child's death. This is to say, unlike what was the case with FSO, an argument from dependence succeeds in showing that I am responsible for failing to save the child in *Planted Sharks*. The argument appeals to the right kind of dependence, namely, causation, and it goes as follows:

- (7) In *Planted Sharks*, I am responsible for planting the sharks yesterday.
- (8) My planting the sharks yesterday caused the child's death today.
- (9) TR (Causal): An agent's responsibility for X transmits to an outcome Y iff X causes Y (and the required provisos are met).
- (10) Therefore, I am responsible for the child's death today.
- (11) If I am responsible for the child's death, then I am responsible for my failure to save the child.
- (12) Therefore, in *Planted Sharks*, I am responsible for my failure to save the child.

This argument seems to go through. I am clearly responsible for planting the sharks, and my planting the sharks is clearly a cause of the child's death. Thus, it follows by TR (Causal) that I am responsible for the child's death. And, if I am responsible for the child's death, I am presumably also responsible for my failure to save the child.<sup>19</sup>

In this section, I presented the *right* reasons for rejecting OA. I argued that *Planted Sharks* is a counterexample to OA, and I used TR (Causal) to support this claim. In the next section, I will argue that a *different* asymmetry thesis is unscathed by *Planted Sharks*. Again, this new asymmetry thesis will be supported, at least in part, by TR (Causal).

## 5. The new moral asymmetry

In the last section I pointed out that, in *Planted Sharks*, the child's death today was caused by my planting the sharks in the water yesterday. Presumably, however, it *wasn't* also caused by my failure to jump into the water today. After all, jumping in today wouldn't have helped. Otherwise put: in *Planted Sharks*, I brought about the child's death today by bringing about *yesterday* that I could not prevent that death today, *not* by not trying

to prevent a death *today* that I couldn't have prevented. In Planted Sharks, then, the fact that the child's death would have occurred anyway deprives my omission to jump in of its causal powers with respect to the death.

Now, this suggests that there might be a new asymmetry between actions and omissions lurking in the background. For actions don't generally behave in this way. Take my action of planting the sharks, for instance. Imagine that, unbeknownst to me, had I not planted the sharks, someone else would have (and there's nothing that I could have done to prevent it). Still, given that I planted the sharks, I caused the child's death. In other words, although the child's death would still have occurred had I not planted the sharks, my planting the sharks caused it. By contrast, as I have pointed out, given that the child's death would still have occurred even if I had jumped in, my omission to jump in *didn't* cause it.

The asymmetry between actions and omissions that seems to be lurking in the background is a *causal* asymmetry: it is an asymmetry in the conditions under which actions and omissions can be causes. It is expressed by the following claim, the *New Asymmetry (Causal)* claim:

NA (Causal): An action can cause an outcome even if the outcome would still have occurred in the absence of the action. By contrast, an omission cannot cause an outcome if the outcome would still have occurred in the absence of the omission.

According to NA (Causal), an agent can cause an outcome by *acting* a certain way even if the outcome would still have occurred had he not acted that way, but an agent cannot cause an outcome by *failing to act* a certain way if the outcome would still have occurred had he acted that way.

Of course, a few examples aren't enough to support a general principle like NA (Causal), and, in particular, the general claim that it makes about omissions. I will argue for NA (Causal) in the next section. Now, however, I will show that, *on the assumption that NA (Causal) is true*, a new moral asymmetry emerges between actions and omissions. This new moral asymmetry is based on NA (Causal) and TR (Causal), the principle of transmission of responsibility that I defended in section 3.

Recall TR (Causal):

TR (Causal): An agent's responsibility for X transmits to an outcome Y iff X causes Y (and the required provisos are met).

From TR (Causal) and NA (Causal), this claim follows:

NA: An agent's responsibility for an action can transmit to an outcome even if the outcome would have occurred anyway in the absence of the action. However, an agent's responsibility for an omission cannot transmit to an outcome if the outcome would have occurred anyway in the absence of the omission.

NA is the *New Asymmetry* claim. It is a moral asymmetry between actions and omissions: an asymmetry in the conditions under which one can be responsible for outcomes *by* action and *by* omission (in virtue of having acted a certain way and in virtue of having failed to act a certain way). Alternatively put, it is an asymmetry in the conditions under which responsibility for actions and omissions *transmits* to outcomes.<sup>20</sup>

Take the example of Sharks, for instance. In Sharks, I am responsible for an *omission* of mine, failing to jump into the water (since, as far as I knew, I could have saved the child by jumping in). Now, given the presence of the sharks, the child's death would still have occurred had I jumped in. Thus NA entails that my responsibility for failing to jump in *doesn't* carry over to the child's death. This is the right result. Contrast this with Planted Sharks. In Planted Sharks, too, the child's death would still have occurred had I jumped in. However, in contrast with Sharks, in Planted Sharks I am responsible for an *action* that I performed yesterday—namely, planting the sharks—in virtue of which I cannot save the child today. Thus, NA allows us to say that I am responsible for the child's death because I am responsible for planting the sharks.<sup>21</sup> Notice that NA allows us to say this even if it turned out that, had I not planted the sharks yesterday, someone else would have, and thus, even if it turned out that, had I not planted the sharks yesterday, the child's death would still have occurred today. Again, this is the right result.<sup>22</sup>

Here are two further scenarios that serve to illustrate the content of NA. In the first scenario, "Drunken Doctor," a tumor is found in a patient's brain, and the surgeon who is to operate on the patient gets drunk before the operation. During the operation, he fails to perform a cut that would have served to remove the tumor, and the patient dies as a result. Had the doctor not been drunk, however, he would have recognized the need for the cut, he would have made the cut and removed the tumor, and the patient would have lived. Clearly, the doctor is responsible for the death of his patient in Drunken Doctor. In the second scenario, "Lucky Doctor," a tumor is found in a patient's brain, and the doctor negligently fails to make the cut that doctors normally make in cases of that sort to try to remove the tumor. The patient dies from the tumor. However (unbeknownst to the doctor at the time of the operation), the tumor was too deep into the patient's brain and couldn't have been removed, so, as it turns out, making the cut wouldn't have helped (imagine that there is no way to remove tumors that deep without killing the patient). In this case, it seems that, although the doctor is responsible for not making the cut (any good surgeon would have made the cut in his place), he is not similarly responsible for the patient's death. In this respect, the doctor is morally lucky: as it turned out, the patient wouldn't have lived if he had made the cut, so he is not responsible for the death by virtue of not making the cut. Thus, Drunken Doctor and Lucky Doctor differ with respect to the doctor's responsibility

for the patient's death: while the drunken doctor is responsible for the death, the lucky doctor is not.<sup>23</sup>

What would NA say about these cases? About Lucky Doctor, NA would say that, even if the doctor is responsible for an omission of his, his not making the cut, the responsibility does not carry over to the patient's death. For the death would still have occurred even if he had made the cut. This is the right result. Now consider Drunken Doctor. In Drunken Doctor, had the doctor made the cut, the patient *wouldn't* have died. Hence NA does *not* entail that the doctor's responsibility for his failure to make the cut doesn't carry over to the death. NA only says that an agent's responsibility for his omission does not carry over to an outcome when the outcome would *still* have occurred in the absence of the agent's omission. In cases like Drunken Doctor, however, the outcome would *not* have occurred in the absence of the agent's omission. Therefore, NA is consistent with the claim that the drunken doctor is responsible for the patient's death in virtue of his not making the cut. Again, this is the right result.<sup>24</sup>

In this section I motivated NA (Causal) by appeal to examples, and I explained how the new moral asymmetry claim, NA, follows from NA (Causal) and TR (Causal). I already argued for TR (Causal) in section 3. Thus, in order to complete my argument for NA, I must argue for NA (Causal). I turn to this in the next section.

## 6. The causal asymmetry

Recall NA (Causal):

NA (Causal): An action can cause an outcome even if the outcome would still have occurred in the absence of the action. By contrast, an omission cannot cause an outcome if the outcome would still have occurred in the absence of the omission.

Why should we believe that there is such a difference between the causal powers of actions and omissions? In what follows, I offer a reason to believe this.

Start by contrasting the following two types of situations, which I will call *switches* and *non-switches*. An example of a *switch* is the following case:

*Redirecting the Train:* A train is hurtling down a track, where a person, Victim, is standing (up ahead). There is a spur of track on the right, and a switch. I flip the switch and as a result the train turns; however, the two tracks reconverge before the spot where Victim is standing. Victim gets killed by the train.

By contrast, here is an example of a *non-switch*:

*Pulling the Trigger:* Assassin is about to shoot Victim. I don't realize this, and since I want Victim to die too, I pull the trigger of my own gun. Victim dies as a result.

Here is yet another *non-switch*:

*Removing the Shield*: Assassin is about to shoot Victim. There is a bulletproof shield that would have stopped the bullet before it reached Victim. I remove the shield. Assassin shoots and Victim dies as a result. Had I not removed the shield, however, Backup would have, and Victim would still have died.

Fill in the details of the cases so that in none of these cases there is something that I could have done to save Victim: Victim would still have died, had I done anything else. (Imagine, in the first case, that I cannot stop the train and that there are no other tracks to which to divert the train; in the second case, that I cannot stop Assassin; and, in the third case, that I cannot stop either Assassin or Backup.) If so, in all of these cases, switches and non-switches, what I do does not affect the outcome—Victim would still have died, had I done anything else—but merely helps to determine the route by which the outcome occurs.

However, intuitively, there is a causal difference between switches and non-switches. Intuitively, I *don't cause* the death in a switch like Redirecting the Train, but I *cause* the death in a non-switch like Pulling the Trigger or Removing the Shield. Intuitively, when one merely redirects a train that was already going to hit and kill a person, one doesn't thereby cause the person's death; by contrast, when one shoots the bullet that kills a person, or when one removes an obstacle to an ongoing bullet, one does cause the person's death, even if the death would still have occurred had one not done those things.<sup>25</sup>

Now, what is the source of this causal difference between the two types of case? Intuitively, it seems to be the following. Even though in both types of case the outcome would still have occurred had I failed to act the way I did, in a *non-switch* the agent creates a new threat or promotes a preexisting threat. Thus, in Pulling the Trigger, the death would have occurred even if I hadn't shot, because Assassin would have shot, but by pulling the trigger I created a new threat that went to completion. Hence, I caused his death. And, in Removing the Shield, the death would have occurred even if I hadn't removed the shield, because Backup would have removed it, but by removing the shield I promoted the old threat to Victim: I removed an obstacle to it, and the threat then went to completion. Hence, I caused Victim's death. By contrast, in a *switch*, there is a preexisting threat and the agent does not promote it or create a new one. At most, the agent diverts the preexisting threat onto a slightly different path, as in Redirecting the Train. Hence, the agent does not cause the outcome by acting the way he does.

I will assume that this account of the causal difference between switches and non-switches is intuitive enough and can be put to work successfully. A more complete explanation of the difference should probably contain an account of the concepts of a threat, creation of a threat, and promotion of a threat. Unfortunately, I cannot do this here. Thus, I have to rely on

intuition to ground my claims about the difference between switches and non-switches. But it does seem intuitive enough, first, that there is a causal difference between switches and non-switches and, second, that the causal difference is grounded in a difference between creating a new threat or promoting an old threat, on the one hand, and neither creating a new threat nor promoting an old threat, but at most redirecting an old threat, on the other.

Now, assuming this story about the causal difference between switches and non-switches, I suggest that a similar story can be told in support of NA (Causal). NA (Causal) says that the fact that an outcome was going to happen anyway deprives an agent's omission of its causal powers towards the outcome, while it doesn't necessarily deprive an agent's action of its causal powers towards the outcome. My suggestion is that this is so because every omission case where the outcome would have occurred anyway is a *switch*, but some action cases where the outcome would have occurred anyway are *non-switches*. We have already seen that some action cases where the outcome would have occurred anyway are non-switches (Pulling the Trigger and Removing the Shield are cases of this type). In order to see that NA (Causal) is true, then, it remains to be shown that any omission case where the outcome would have occurred anyway is a switch. I turn to this now.

Consider Sharks again. Given the presence of the sharks in the water, I couldn't have prevented the death of the child by jumping in. Also, intuitively, by refraining from jumping in, I didn't create a new threat or promote the preexisting threat in any way, simply because I didn't *do* anything: I just let the ongoing process run its preset course. There was an existing process leading to the death, which was developing in a certain way, and, by failing to act, I let it keep developing that way. Intuitively, any time one doesn't interfere with a process, one doesn't promote it (or create a new one): one just lets it unfold. Thus, it seems that omissions don't promote existing threats or create new threats; at most, they let existing threats develop in certain ways. If so, any omission case where the outcome would have occurred anyway is a switch, in the way I have defined switches.

A word of clarification: I stress that I am *not* suggesting that, given that omissions don't promote existing threats or create new threats, they never cause anything. Many times, letting a process develop in a certain way *is* sufficient for causing an outcome. For instance, if a mother doesn't feed her baby when she could easily have fed him, and the baby starves to death, then the mother doesn't just let the process of starvation develop, but, in addition, she causes the baby's death. Rather, what I am suggesting is that, *when the outcome would have occurred anyway*, an omission doesn't cause the outcome given that, first, the outcome would have occurred anyway and, second, omissions don't promote or create threats. This is what happens in Redirecting the Train, an ordinary switch: in Redirecting the Train,

I don't cause Victim's death by flipping the switch because, first, the death would still have occurred if I hadn't flipped the switch, and, second, by flipping the switch I don't promote the preexisting threat or create a new one.

Also, it is important to realize that, although NA (Causal) says that, when an outcome would still have occurred in the absence of an omission, *that omission* isn't a cause of the outcome, this is consistent with its being the case that something else that the agent did or failed to do caused the outcome. For instance, in *Planted Sharks*, my failure to jump in isn't a cause of the child's death, because the death would still have occurred if I had jumped in, but *my planting the sharks*, an action of mine, is a cause. Or imagine that I am an absentminded lifeguard. The child is drowning, and there are no sharks. Still, I wouldn't have saved the child by jumping in because, if I had jumped in, I would have forgotten the life preserver and thus I wouldn't have been able to save the child. According to NA (Causal), my failure to jump in isn't a cause of the child's death, because the death would still have occurred if I had jumped in (since I would have jumped in without the life preserver). However, arguably, *my failure to jump in with the life preserver*, a *different* omission of mine,<sup>26</sup> is a cause of the child's death.

Now, all the cases that I discussed have a common structure. Their structure is the following. An agent fails to act in a certain way and an outcome occurs. Had the agent acted in the relevant way, the outcome would still have occurred because the agent's intervention wouldn't have sufficed to stop the ongoing threat, which would have still gone to completion. For instance, in *Sharks*, the drowning process is on its way, and I wouldn't have stopped it by jumping into the water: it would have gone to completion anyway. Call any case of this type an *easy* case. I think that the preceding discussion succeeds in showing that, at least in the easy cases, the agent's omission doesn't cause the outcome.<sup>27</sup> However, there are cases of a different type, with a different underlying structure, which are more likely to make trouble for NA (Causal). I turn to these cases in the next section.

### **7. The hard cases and the issue of transitivity**

A *hard* case is a case with the following structure. An agent fails to act in a certain way and an outcome occurs. Had the agent acted in the relevant way, the outcome would still have occurred; however, it would have occurred as a result of a *different* threat from the one that actually led to the outcome (call the threat that actually led to the outcome "the *actual* threat"). That is, the agent could have stopped the actual threat, but he still couldn't have prevented the outcome. He couldn't have prevented the outcome because there was a backup threat: had he stopped the actual threat, then the backup threat would have issued in the outcome just the same. Here is an example:

*Evil Bystander*: Assassin is about to shoot Victim. I could stop him, but, since I want Victim dead too, I let him shoot. Victim dies. Unbeknownst to me, however, Backup is waiting in reserve. Had I stopped Assassin, then Backup would have shot, and Victim would still have died (imagine that I couldn't have stopped Backup, and imagine, also, that the details of the death would have been very similar, so it would have been the same death).

In *Evil Bystander*, I could have stopped the actual threat, but I still couldn't have prevented the outcome from happening because, had I acted, the backup threat would have issued in the outcome just the same.

Why is *Evil Bystander* a *hard* case for NA (Causal)? Because, in this case, one might be tempted to say the following. Given that I could have stopped the process that actually ended in Victim's death and I didn't, then, by failing to stop it, I contributed to the death. So my failure to stop Assassin was one of the causes of the death, although I couldn't have prevented the death. Hence, NA (Causal) is false: an omission can cause something that would have occurred anyway. In other words, the objection to which a hard case like *Evil Bystander* gives rise is that, contrary to what I have suggested, merely letting a threat unfold can be enough to cause an outcome, even if the outcome would have occurred anyway. In particular, the objection is that letting the *actual* threat unfold, when one could have stopped it, is enough to cause an outcome, even if the outcome would have occurred anyway via a backup threat.<sup>28</sup>

In what follows, I argue that whether we think that we should revise NA (Causal) in light of hard cases like *Evil Bystander* essentially depends on our take on the transitivity of causation: we should revise it if we think that causation is transitive, but not so if we think that it is not transitive, or if we are unsure. I will briefly explain the main reasons for being skeptical about transitivity. However, I will not try to settle the issue of transitivity here, because it is an issue that deserves a much more extensive treatment than I can offer in this paper. As a result, I will leave it open whether NA (Causal) is in need of refinement, and I will briefly indicate how NA (Causal) *could* be revised in light of the hard cases, if one were to deem it necessary.

Let us start by asking: why should we believe that, if one lets the actual threat unfold when one could have stopped it, one causes the outcome? The main motivation for believing this is the following. Since one could have stopped the actual threat but doesn't, one's failure to act is a cause of the actual threat's unfolding,<sup>29</sup> and the actual threat's unfolding is a cause of the outcome; hence, by transitivity, one's failure to act is a cause of the outcome. In *Evil Bystander*, my failure to stop Assassin was a cause of his shooting, and Assassin's shooting was a cause of Victim's death; hence, by transitivity, my failure to stop Assassin was a cause of Victim's death. In other words, the thought that, in the hard cases, the agent's omission causes the outcome because it lets the actual threat go to completion, is fueled by the thought that causation complies with the following principle:

*Transitivity*: If X causes Y, and Y causes Z, then X causes Z.

In the contemporary literature on causation, there is considerable debate over whether Transitivity is true. Some people believe that Transitivity is a general truth about causation, but other people think that there are counterexamples.<sup>30</sup> As I said, I will not try to settle this issue here; however, I will briefly explain what some of the problems with endorsing Transitivity are.

One of the problems with Transitivity is illustrated by Evil Bystander itself. As I have pointed out, Transitivity supports the claim that my failure to stop Assassin causes Victim's death in Evil Bystander: if Transitivity is true, then my failure to stop Assassin causes Victim's death because it causes Assassin's shooting, which causes Victim's death. However, by the same token, Transitivity supports the claim that, *had I stopped Assassin*, then *my stopping Assassin* would have caused Victim's death. For, had I stopped Assassin, then my stopping Assassin would have caused Backup to shoot, which would have caused Victim's death (recall that Backup is determined to shoot if Assassin doesn't). Hence, if Transitivity were true, it would follow that I would have also caused the death if I had stopped Assassin. And this is counterintuitive.<sup>31</sup>

Transitivity also yields puzzling results about Sharks. As we have seen, my failure to jump into the water in Sharks is not a cause of the child's death. But imagine that I *had* jumped in. Had I jumped in, then my jumping in would have caused the sharks to attack me, and the shark attack would have in turn caused the child's death.<sup>32</sup> Hence, if Transitivity were true, it would follow that my jumping into the water in Sharks would have caused the child's death. In other words, I *don't* cause the child's death by refraining from jumping in, but, if Transitivity is true, I *would* have caused his death by jumping in. This is an extremely counterintuitive result.

I have explained how our take on the hard cases seems to heavily depend on our take on Transitivity, and I have explained what some of the main problems with Transitivity are. Now, could it be argued that the hard cases pose a problem for NA (Causal) *regardless* of what our take on Transitivity is? In particular, could it be argued that *intuition* alone shows that the hard cases are counterexamples to NA (Causal)?

I will argue that intuition alone is not enough to show this, because our intuitions about the hard cases are morally tainted. Take Evil Bystander. The way I have set up the case, Evil Bystander is a case where my intentions are morally despicable. For I failed to stop Assassin when I thought that I could have prevented Victim's death by stopping him (since I was unaware of Backup's intentions). No wonder, then, that we feel tempted to say that I caused Victim's death. However, imagine that the details of the case are different. Imagine, for instance, that I am a good person and that I don't want for Victim to die. However, I still fail to stop Assassin. I fail to stop him because I see that Backup is waiting in reserve and I know that he will

shoot if Assassin doesn't. Imagine, moreover, that I have good reason to let Assassin shoot instead of Backup: imagine, for instance, that I know that Backup is a very good shooter and I think that Victim will end up suffering more if I stop Assassin (say, because his hopes will be high again for a moment and then there will come the tragic realization that he will die anyway). In such circumstances, I find it a lot less plausible to say that I cause Victim's death by not stopping Assassin. After all, all I did was to do nothing, when there was nothing that I could have done to help.<sup>33</sup>

If so, this should be enough to cast doubt on our intuitions about Evil Bystander itself, where my intentions were bad. For none of the details of the sort I mentioned should be relevant to whether I caused Victim's death by failing to stop Assassin, since they only have to do with my intentions and beliefs (with what I thought was the case and with what I intended to be the case, but not with what was actually the case). Such details help to determine whether I did the morally right thing in failing to stop Assassin, but not whether I caused Victim's death by failing to stop Assassin. This suggests that our intuitions about these cases are tainted by our moral judgments about the cases. Hence, we should not trust those intuitions.

I argued that, apart from Transitivity, there doesn't seem to be any compelling reason to believe that the hard cases undermine NA (Causal). Hence, unless we are moved by Transitivity, we shouldn't revise NA (Causal) in light of the hard cases. I also explained what some of the problems with Transitivity are. However, for all I have said, there might be powerful reasons to hold Transitivity despite some of its counterintuitive consequences. If so, NA (Causal) would be in need of refinement. In light of this, I will end my discussion of the causal asymmetry with a sketch of how NA (Causal) could be revised, if it were necessary to revise it.

If one were to revise NA (Causal) in light of the hard cases, then one should most likely proceed by restricting it to the conditions described in the *easy* cases. The easy cases, remember, are cases where the actual threat would have inevitably gone to completion, regardless of what the agent did. Thus, roughly, the revised causal asymmetry would be the claim that, *in circumstances where the actual process cannot be stopped*, omissions cannot be causes but actions still can. Sharks illustrates the claim about omissions: given the presence of the sharks, my jumping in wouldn't have stopped the ongoing drowning process, which would have gone to completion anyway; hence, my failure to jump in isn't a cause of the child's death. In turn, Removing the Shield (the example from the last section) illustrates the claim about actions: given the presence of Backup (who would have removed the shield had I not done so myself), my refraining from removing the shield wouldn't have stopped the ongoing shooting process started by Assassin, which would have gone to completion anyway; however, my removing the shield is a cause of Victim's death.

In other words, the revised causal asymmetry would be the claim that, when the actual threat cannot be stopped, an action can still cause the outcome because it can promote the threat, but an omission cannot cause the outcome because it cannot promote the threat. If one were to revise the causal asymmetry this way, one would then have to revise the *moral* asymmetry accordingly. Roughly, the revised moral asymmetry would be the claim that, when the actual threat cannot be stopped, the responsibility for an omission cannot transmit to the outcome but the responsibility for an action can.

Let me sum up the results of the last two sections. I distinguished the easy cases from the hard cases. I argued that the easy cases clearly support NA (Causal). Then I explained why it might be thought that the hard cases make trouble for NA (Causal). Although I expressed some reservations about the hard cases, I argued that, ultimately, the decision whether to revise NA (Causal) in light of those cases is likely to depend on the answer to the question whether causation is transitive. Finally, I explained how NA (Causal) could be revised if causation were transitive.

## 8. Conclusions

In this paper I argued for a new moral asymmetry between actions and omissions. I first explained my reasons for rejecting the old asymmetry thesis that some philosophers have suggested, and I then argued for a new way of understanding the moral asymmetry that the old asymmetry thesis attempted to rescue. On this new way of understanding the moral asymmetry, there is a moral asymmetry between actions and omissions because there is a more fundamental and purely *causal* asymmetry between actions and omissions, and because causation plays a fundamental role in the transmission of responsibility.

## Notes

\* Special thanks to Juan Comesaña, Ned Hall, Judith Thomson, and Stephen Yablo. Many thanks also to Eduardo Flichman, Elizabeth Harman, Sarah McGrath, and audiences at MIT, NYU, University of Michigan-Ann Arbor, SUNY-Buffalo, University of Missouri-Columbia, and the 2003 Pacific APA meeting, where previous versions of this paper were presented.

<sup>1</sup> A case by J. M. Fischer and M. Ravizza (Fischer and Ravizza (1998), p. 125).

<sup>2</sup> After H. Frankfurt (Frankfurt (1969)).

<sup>3</sup> If you are not convinced, imagine that, in addition to the sharks, there is an impenetrable wall of rocks in the water, and a big net, and all kinds of obstacles, each one of which would have prevented me from saving the child; it would be very implausible to say that I am responsible for the child's death in those circumstances. This suggests that our responsibility for outcomes is partly subject to luck (see Nagel (1979)).

<sup>4</sup> See, e.g., van Inwagen (1978). OA's claim about actions gives rise to questions that I will have to set aside here. For instance, does Shooting really show that I can be responsible for an action that I couldn't have failed to perform? If the neuroscientist had intervened by forcing me

to decide to pull the trigger, would my behavior have been an *action* of mine? (Would I have been the *agent* of that behavior?)

<sup>5</sup> In particular, my proposal will not be that omissions can never make us responsible for anything because they can never cause anything (for an example of this view, see Weinryb (1980)). One reason to allow for the existence of causation by omission is that, arguably, if we didn't allow for it, it would follow that there is also a lot less "positive" causation than we normally think there is (for an argument for this view, see Schaffer (2000)).

<sup>6</sup> See, e.g., Clarke (1994) and McIntyre (1994).

<sup>7</sup> Both Clarke and McIntyre suggest an argument along these lines in Clarke (1994) and McIntyre (1994).

<sup>8</sup> This is also the natural way to look at cases in which I am *not* responsible. For instance, as I pointed out in section 1, it is natural to think that, in Sharks, I am not responsible for my failure to save the child because I am not responsible for the child's death.

<sup>9</sup> Recall that I am assuming that omissions can be causes. But, what if the Frankfurt-style objector held a much *narrower* conception of causation according to which omissions cannot be causes? Or, what if he held a much *more permissive* conception of causation according to which my failure to decide to jump in is a cause, but my decision not to jump in is *also* a cause? For reasons that will become apparent in the next section, if he said any of these two things, then it would follow that, contrary to what we are assuming, causation *isn't* appropriate to transmit responsibility, and thus a different premise of the argument, (3'), would fail.

<sup>10</sup> Notice that I am rejecting the view that some philosophers have taken, according to which omissions are a subclass of actions. I do not find this view compelling. One serious problem with it is that sometimes, given an omission, there aren't *any* actions that can plausibly be identified with it, or, alternatively, there isn't a *single* action that can plausibly be identified with it. For discussion, see Weinryb (1980).

<sup>11</sup> Given that everybody agrees about Sharks, everybody should agree about this case. In particular, the Frankfurt-style objector to OA should. Recall that he wants to draw a difference between FSOC and Sharks. He wants to say that, if I couldn't have made the morally right decision (as in FSOC), then I am responsible for failing to save the child; by contrast, if I could have made the morally right decision but I would have been stopped *afterwards* (as in Sharks), then I am not. The train case is analogous to Sharks, not to FSOC: by assumption, the neuroscientist would have intervened *after* I made the decision to switch to C, by forcing me to switch to B.

<sup>12</sup> Notice that it was foreseeable that my flipping the switch to B would be followed by Victim's death. Hence, the failure of transmission of responsibility cannot be blamed on the required epistemic provisos, which are met in this case.

<sup>13</sup> *Pace* some theories of causation that regard counterfactual dependence as sufficient for causation. See, e.g., Lewis (1986).

<sup>14</sup> Notice the parallel with my discussion of FSOC in the preceding section.

<sup>15</sup> Again, notice the parallel with my discussion of FSOC. In FSOC, *had* I been responsible for failing to decide to jump in, I would *also* have been responsible for the child's death, because my failing to decide to jump in caused the child's death. However, as I pointed out in the last section, it is not clear that I was responsible for failing to decide to jump in (what *is* clear is that I was responsible for *deciding* not to jump in).

<sup>16</sup> J. Feinberg endorses a view of this type in Feinberg (1970).

<sup>17</sup> See also Clarke's drunk driver example in Clarke (1994).

<sup>18</sup> Couldn't I have saved him by not planting the sharks yesterday? To dissipate any worries of this sort, imagine that someone else would have planted the sharks had I not done so myself. Then I couldn't have saved the child, not even by not planting the sharks.

<sup>19</sup> Two words of clarification. First, notice that TR (Causal)'s required provisos are met in Planted Sharks. By assumption, I was guilty of negligence in planting the sharks. So, in particular, it was foreseeable that my planting the sharks was likely to be followed by a harm of the sort that ensued. Second, step (11) also seems warranted. For start by imagining that

today there are *other* people around the child that cannot save him as a result of my having planted the sharks yesterday. Clearly, I would be responsible for these other people's failures to save the child in that case. Similarly, if I am the only person standing close to the child today and if I negligently released the sharks yesterday, I am responsible for *my own* failure to save the child today. (Thanks to Iris Einheuser on this last point.)

<sup>20</sup> It is important to distinguish NA from a different asymmetry that some people have defended, according to which there is a moral difference between causing a death by action and causing it by omission, or between killing someone and letting someone die. Some people have suggested that, other things being equal, killing is worse than letting die. But the discussion of whether this is true centers around cases where the outcome *wouldn't* have occurred anyway, in particular, it is generally assumed that one *lets* someone die by failing to act only if, by acting, one *would* have saved the person. Thus, the question whether killing is worse than letting die is not likely to shed any light on the question whether NA is true. For discussion of the killing and letting die distinction, see, e.g., Bennett (1994) and Foot (1994).

<sup>21</sup> That is, I am responsible for the child's death, but the responsibility is inherited from my planting the sharks, not from my failing to jump in. It is inherited from the former and not from the latter because the former, not the latter, is a cause of the death.

<sup>22</sup> NA is only *consistent* with this claim, instead of *entailing* this claim, because—unlike TR (Causal)—it doesn't provide *sufficient* conditions for the transmission of responsibility (it only claims that a certain condition is *necessary* for the transmission of responsibility in the case of omissions but not in the case of actions).

<sup>23</sup> I emphasize that we would probably still want to say that both doctors are *bad* doctors, maybe even *equally bad* doctors.

<sup>24</sup> Again, NA is only *consistent* with this claim instead of *entailing* it because it doesn't provide sufficient conditions for the transmission of responsibility (see n. 22).

<sup>25</sup> I stress that this is only an appeal to intuition. Some theories of causation don't yield the intuitive verdict about switches. For attempts at capturing our causal intuition about switches, see Rowe (1989) and Yablo (2002).

<sup>26</sup> Presumably, my failure to jump in and my failure to jump in with the life preserver are different failures, since they obtain under different circumstances.

<sup>27</sup> There is a particular type of easy case for which this might be a bit more controversial. Imagine that *two* agents had to intervene for a threat to be stopped and that each independently refrains (suppose, e.g., that two buttons had to be pushed for an explosion to be prevented and that each agent fails to push one). In my (2004), I argue that the collective omission is a cause in this case, but that, as NA (Causal) predicts, the agents' individual omissions aren't. However, I then argue that we can still say that each of the agents is responsible for the outcome (in a way that is consistent with the idea that causation is required for the transmission of responsibility): each is responsible for the outcome because each is responsible for the collective omission, which causes the outcome.

<sup>28</sup> J. J. Thomson believes that omissions can cause outcomes that would have occurred anyway in cases of this sort (Thomson (2003), p. 81).

<sup>29</sup> This is the difference from the *easy* cases. In an easy case, I don't cause the actual threat's unfolding, since I couldn't have stopped the actual threat.

<sup>30</sup> For discussion of Transitivity, see, e.g., Hall (2000), Hitchcock (2001), Paul (2000), and my (2005).

<sup>31</sup> In my (2005) I provide an argument against the existence of this type of causal structure.

<sup>32</sup> Just as, if I attack someone that is about to rescue the child, and then the child dies, I am a cause of the child's death.

<sup>33</sup> I don't mean to suggest that it is impossible to say that I cause the death in this case (clearly, it would be possible to say that I cause the death but I am still not morally responsible for it). All I mean to suggest is that our intuitions about this case are different from those about Evil Bystander.

## References

- Bennett, J. (1994) "Whatever the Consequences," reprinted in B. Steinbrock and A. Norcross (eds.), *Killing and Letting Die*, New York: Fordham University Press, pp. 167–91.
- Clarke, R. (1994) "Ability and Responsibility for Omissions," *Philosophical Studies* 73, pp. 195–208.
- Feinberg, J. (1970) "Sua Culpa," in *Doing and Deserving: Essays in the Theory of Responsibility*, Princeton: Princeton University Press, pp. 187–221.
- Fischer, J. M. and M. Ravizza (1998) *Responsibility and Control*, Cambridge: Cambridge University Press.
- Foot, P. (1994) "Killing and Letting Die," reprinted in B. Steinbrock and A. Norcross (eds.), *Killing and Letting Die*, New York: Fordham University Press, pp. 280–89.
- Frankfurt, H. (1969), "Alternate Possibilities and Moral Responsibility," *Journal of Philosophy* 66, pp. 829–39.
- Hall, N. (2000), "Causation and the Price of Transitivity," *Journal of Philosophy* 97, 4, pp. 198–222.
- Hitchcock, C. (2001) "The Intransitivity of Causation Revealed in Equations and Graphs," *Journal of Philosophy* 98, 6, pp. 273–99.
- Lewis, D. (1986) "Causation," reprinted in *Philosophical Papers*, vol. II, New York: Oxford University Press, pp. 159–213.
- McIntyre, A. (1994) "Compatibilists Could Have Done Otherwise," *Philosophical Review* 103, 3, pp. 458–88.
- Nagel, T. (1979) *Mortal Questions*, New York: Cambridge University Press.
- Paul, L. (2000) "Aspect Causation," *Journal of Philosophy* 97, 4, pp. 235–56.
- Rowe, W. (1989) "Causing and Being Responsible for What is Inevitable," *American Philosophical Quarterly* 26, 2, pp. 153–9.
- Sartorio, C. (2004), "How to be Responsible for Something without Causing It," *Philosophical Perspectives*, Ethics, v. 18, 1, pp. 315–336.
- Sartorio, C. (2005), "Causes as Difference-Makers," forthcoming in *Philosophical Studies* 123, 1–2, pp. 71–96.
- Schaffer, J. (2000) "Causation by Disconnection," *Philosophy of Science* 67, pp. 285–300.
- Thomson, J. J. (2003) "Causation: Omissions," *Philosophy and Phenomenological Research* 66, 1, pp. 81–103.
- van Inwagen, P. (1978) "Ability and Responsibility," *Philosophical Review* 87, pp. 201–24.
- Weinryb, E. (1980) "Omissions and Responsibility," *The Philosophical Quarterly* 30, 118, pp. 1–18.
- Yablo, S. (2002) "De Facto Dependence," *Journal of Philosophy*, pp. 130–48.