

Failing to Do the Impossible*

1. The billionaire puzzle

A billionaire tells you: “That chair is in my way; I don’t feel like moving it myself, but if you push it out of my way I’ll give you \$100.” You decide you don’t want the billionaire’s money and you’d rather have him go through the trouble of moving the chair himself, so you graciously turn down the offer and go home. As it turns out, the billionaire is also a stingy old miser; he was never willing to let go of \$100. Knowing full well that the chair couldn’t be moved due to the fact that it was glued to the ground, he simply wanted to have a laugh at your expense.

This is a case of omission: you expressed your agency mainly by what you intentionally omitted to do, not by what you did. But what exactly did you intentionally omit to do in this case? There is no question that you intentionally omitted to *try* to move the chair. But did you also intentionally omit to *move the chair*? On the face of it, it seems that you didn’t. Even if not moving the chair was something that you wanted to do, and even if, as a matter of fact, you didn’t move the chair, it seems that you didn’t intentionally omit to move the chair. For, contrary to what you were led to believe, you couldn’t move it, and the fact that you couldn’t move it seems to preclude your having intentionally omitted to move it. (Of course, this is not to say that you shouldn’t be commended for failing to comply with the billionaire’s whimsical wants. You should still be commended for that, but not by virtue of having intentionally omitted to move the chair; only by virtue of having intentionally omitted to try to move it. After all, that’s what he wanted you to do: he wanted you to try to move the chair motivated by the desire to get \$100.)

Note that the claim is that you didn't *intentionally omit* to move the chair. Perhaps there is some ordinary or theoretically useful sense of the word "omission" in which there is an omission here (and, in fact, sometimes I will talk as if this were true). Still, what's important for our purposes is that this isn't an intentional omission. For example, if omissions were mere non-doings, then, clearly, there would be an omission even in this case. If omissions were mere non-doings, it would be very easy to omit to do the impossible: all of us would always omit to do the impossible, since it's impossible. But there would still be restrictions on what we can intentionally omit to do. In other words, even if omitting to do the impossible were easy (as easy as it can get), intentionally omitting to do the impossible would still be hard, and for what is apparently the same reason: because the act in question is impossible.

Examples like the billionaire case suggest that intentional omissions can be importantly subject to *counterfactual constraints*, i.e. constraints that have to do with what the agent could have done, or with what the agent does in other possible worlds different from the actual world. Moreover, these seem to be constraints that only omissions are subject to (unlike actions of the positive kind) or, at the very least, that mainly omissions are subject to. For (positive) actions don't seem to be subject to the same kind of constraints: on the face of it, whether an agent acts intentionally is wholly dependent on the actual facts; it is not determined by what happens in other possible worlds. Compare: what prevents me from intentionally jumping to the moon if I try to jump to the moon is whatever actually stops me when I try (say, the gravitational field). By contrast, what prevents me from intentionally failing to jump to the moon if I intend not to jump to the moon is the fact that I wouldn't have been able to jump to the moon even if I had tried. In other words: I cannot intentionally jump

to the moon, and I cannot intentionally fail to jump to the moon either. But the reason I cannot do the latter is grounded in the counterfactual facts in a way that the reason I cannot do the former is not.

Now, what exactly is the role of counterfactual constraints in intentional omissions? For some theories, that role is straightforward: an agent intentionally omits to act *only if* it was possible for him to perform the act in question (see, e.g., Zimmermann (1981), Ginet (2004), and Bach (2010)). In other words, according to these theories, intentional omissions require “the ability to do otherwise” (or “alternative possibilities”). But, is the ability to do otherwise truly a necessary condition of intentional omissions? Is it really impossible to omit (intentionally) to do the impossible?

Interestingly, other examples suggest otherwise: they suggest that it *is* possible for agents to omit (intentionally) to do the impossible. Consider the following scenario:

The Child and the Neuroscientist: You see a child drowning in the water, you deliberate about whether to jump in to save him and, as a result of your own free deliberation, you decide not to do so. Unbeknownst to you, there was an evil neuroscientist closely monitoring your brain. Had you shown any signs that you were about to decide to jump in, the neuroscientist would have intervened by making you decide not to jump in and, as a result, you would still have failed to rescue the child.

This is a “Frankfurt-style” scenario. In a Frankfurt-style scenario, an agent makes a decision to act or not act completely on his own, and then acts or doesn’t act based on that decision. Unbeknownst to the agent, the evil neuroscientist has been waiting in the wings to make sure

that he makes the relevant decision but, given that the agent arrives at it on his own, the neuroscientist never has to intervene. Frankfurt famously used these scenarios to show that moral responsibility for actions or omissions doesn't require the ability to do otherwise (the claim about actions appears in Frankfurt (1969); he then extended the claim to omissions in (Frankfurt (1994)), as a reaction to a paper by Fischer and Ravizza (Fischer and Ravizza (1991)). In *The Child and the Neuroscientist*, Frankfurt would argue, you are morally responsible for your not saving the child despite the fact that you couldn't have saved him (you couldn't have saved the child, in this case, because you couldn't even have decided to save him). It is equally clear (if not clearer), Frankfurt would presumably say, that you intentionally omit to save the child in this case. You intentionally omit to save to him because you failed to save him on the basis of your own decision not to save him. Again, this is so even if you couldn't have saved him.¹

In other words, *The Child and the Neuroscientist* seems to be a counterexample to *both* of the following principles:

PPA ("Principle of Possible Action" (for moral responsibility), van Inwagen (1978)):

An agent is morally responsible for omitting to perform a given act only if he could have performed that act.

PPA-IO ("Principle of Possible Action for Intentional Omissions"): An agent intentionally omits to perform a given act only if he could have performed that act.

So, this is how things stand. On the one hand, in the billionaire case, we want to say that the agent doesn't intentionally omit to act because he couldn't have done otherwise. However, on the other hand, it seems that we cannot say this because, as The Child and the Neuroscientist scenario suggests, agents can intentionally omit to act even when they couldn't have done otherwise. Hence there is a puzzle. I will call this puzzle "the billionaire puzzle."

How should we try to solve the puzzle? One way to solve it would be to endorse a form of skepticism about Frankfurt-style cases. My own view is that there are good reasons to resist the argument against PPA based on Frankfurt-style cases, and that, similarly, there are good reasons to resist the parallel argument against PPA-IO.² Still, in this paper I will assume that Frankfurt is right in believing that Frankfurt-style omission cases are counterexamples to both of these principles. I will do this for two reasons. First, there is no denying that Frankfurt-style cases have intuitive appeal, and that they show that principles like PPA and PPA-IO are, at least, debatable. In fact, most people who have contributed to the literature on moral responsibility for omissions seem to agree with Frankfurt about PPA (even Fischer and Ravizza, who changed their minds after their (1991) paper; see their (1998), ch. 5).³ And I take it that most people would also agree that PPA-IO fails, for similar reasons. So the question arises: what *could* the connection between intentional omissions and alternative possibilities *be*, if it's not that intentionally omitting to act requires alternative possibilities? There surely seems to be some connection (otherwise, why did it seem so plausible to say, in the billionaire case, that the reason you didn't intentionally omit to move the chair is that you couldn't have moved the chair?). But what is this connection, if it's not that you cannot intentionally omit to do what you couldn't have done? Secondly, even if intentional omissions did require alternative possibilities, there would still be some explaining to do as to *why* it is

that they do. Why is PPA-IO true, if it is true? This seems to call for an explanation (in particular, an explanation would be most pressing if it turned out that omissions behave in this way but actions don't—why is there such an asymmetry?). As we will see, reflecting about what the relation between intentional omissions and alternative possibilities could be if it's not PPA-IO would also help us find such an explanation.

Hence the billionaire puzzle is the puzzle that arises if one believes that PPA-IO is false, or if one is unsure whether it is true. The challenge is to explain the sense in which your inability to move the chair in the billionaire case is accountable for your not having intentionally omitted to move the chair, without appealing to the claim that intentional omissions require alternative possibilities. My main aim in this paper is to solve the billionaire puzzle. Of course, there is a similar puzzle that arises about moral responsibility: What is the connection between moral responsibility for omissions and alternative possibilities, if it's not that moral responsibility for an omission requires alternative possibilities? Although this question is not my main concern here, we will touch on it too. In fact, my strategy will be to examine the relationship between intentional omissions and alternative possibilities in light of the debate about moral responsibility and alternative possibilities. As we will see, that debate can shed important light on our topic. I turn to this in the next section.

2. Frankfurt's revised principle of alternative possibilities

When Frankfurt rejected the principle that an agent's moral responsibility for an action requires the agent's ability to do otherwise (the "Principle of Alternate Possibilities," or PAP), he put forth a different principle about the connection between moral responsibility and alternative possibilities (one that he considered to be friendly to a compatibilist view about

determinism and the freedom of the will). Although this principle was originally intended as a claim about (positive) actions, a question that naturally arises given our topic in this paper is whether it could apply to omissions in the same way. In this section I will first argue that it doesn't: that Frankfurt's principle fails to capture the relationship between moral responsibility and alternative possibilities in cases of omission. However, even if the principle fails for omissions, it will be enlightening for our purposes to see *how* it fails. As we will see, looking at the way in which Frankfurt's principle fails will give us the key to solving the billionaire puzzle.

Frankfurt's revised principle about the relationship between moral responsibility and alternative possibilities (for actions) says:

PAP-revised (revised version of the Principle of Alternate Possibilities, Frankfurt (1969)): An agent is not morally responsible for performing an action if he did it *only because* he could not have done otherwise.

For example, Frankfurt argues, when you are coerced into doing something you are not morally responsible for your act to the extent that you did it only because of the (irresistible) threat that was made. If the threat, and only the threat, moved you to act, then you are not morally responsible for your act. This case should be contrasted with a Frankfurt-style case. In a Frankfurt-style case, what actually causes the agent to act are the agent's own reasons. So, Frankfurt thinks, even if the agent couldn't have done otherwise, he is morally responsible for his act in that case, because he acted for his own reasons.

Let's call the factors in virtue of which an agent couldn't have done otherwise *inevitability factors*. Frankfurt's idea is that inevitability factors can only reduce the agent's moral responsibility when they actually move the agent to act (and when they are the only thing that moves the agent to act, in particular, when the agent doesn't act for his own reasons). According to Frankfurt, then, PAP-revised captures the truth, and the *only truth*, behind the idea that moral responsibility for acts and alternative possibilities are linked. As Frankfurt-style cases show, one can be morally responsible for acting even if one couldn't have done otherwise. But when what actually moves one to act are the inevitability factors, and only those factors, one is not morally responsible. Thus, whereas PAP is false, PAP-revised is true: moral responsibility for an act doesn't require having the ability to do otherwise, but it does require not acting as a result of (only) the inevitability factors.

Let's assume that Frankfurt is right and PPA-revised is the true principle about the relationship between moral responsibility and alternative possibilities in cases of action. Could it also be the true principle about such relationship in cases of omission? This is what one would expect if, as Frankfurt seems to believe, the conditions of moral responsibility for actions and omissions (or at least those that have to do with the kind of "control" the agent must have to be morally responsible) are perfectly symmetrical.⁴ The revised principle for omissions would read:

PPA-revised (revised version of the Principle of Possible Action): An agent is not morally responsible for omitting to perform an act if he omitted to perform it *only because* he could not have done otherwise, i.e. only because he could not have performed the relevant act.

Now, what does it mean to say that an agent omitted to perform an act *only because he couldn't have performed the relevant act*? By analogy with the case of actions, we should take it to mean that the agent omitted to act moved only by inevitability factors—those factors in virtue of which he couldn't have performed the relevant act. In other words, the inevitability factors, and those factors only, caused him to omit to act. So we should take PPA-revised to be the claim that an agent is not morally responsible for omitting to perform an act if he omitted to perform the act moved only by inevitability factors, i.e. if inevitability factors, and those factors only, caused him to omit to act.

Note that here we are assuming that omissions can be caused just like actions can, in particular, we are assuming that an agent's omitting to act can be the causal output of different things, including inevitability factors. I will go along with that assumption (although, as will be apparent later, I think that such an assumption is not ultimately essential to the views put forth in this paper). And recall: Frankfurt's claim would be that PPA-revised is the only truth, or the whole truth, behind the idea that responsibility for omissions and alternative possibilities are connected: an agent is not morally responsible for an omission when he is caused to omit to act by inevitability factors, and this is the only significant connection between moral responsibility for omissions and alternative possibilities.

Now, *is* PPA-revised the truth, and the complete truth, behind the idea that responsibility for omissions and alternative possibilities are connected? Although Frankfurt in fact seems to think so⁵, I'll argue that it isn't. Consider an example where the agent's inability to perform the relevant act seems to relieve him of moral responsibility for his omission. Here

is a paradigm example from the literature (Fischer and Ravizza (1991), p. 261, and Fischer and Ravizza (1998), p. 125):

Sharks: You see a child drowning in a pond and you decide against jumping in to save him. Unbeknownst to you, you couldn't have saved the child: if you had jumped into the water, some hungry sharks would have attacked you and prevented you from saving him.

As Fischer and Ravizza claim, you are not morally responsible for not saving the child in this case (only for not trying to save him). Frankfurt agrees (Frankfurt (1994), p. 622). Moreover, the fact that you are not morally responsible for your not saving the child seems to have something to do with your inability to save the child (in fact, intuitively we would say that you are not morally responsible for not saving him because you couldn't have saved him). But in this case you were *not* moved to omit to act by the inevitability factors. The inevitability factors in this case are the sharks. Clearly, the sharks didn't move you to omit to act in this case. You weren't even aware of their presence!⁶

It is not completely clear what Frankfurt would say about this. There is one point in his paper that suggests that he might want to insist that the sharks are operative in the relevant sense. He writes:

“The sharks operate both in the actual and in the alternative sequences, and they see to it that the child drowns no matter what John [the agent in the situation he describes] does.” (Frankfurt (1994), p. 623.)

The sharks certainly operate in the sense that they *guarantee* that the agent will not save the child. But so does the neuroscientist in a Frankfurt-style case, where (we are assuming) the presence of the neuroscientist bears no relevance to the agent's moral responsibility. So this cannot be the relevant sense of "operate": for the sharks to operate in the relevant sense, they must *cause* the agent's omission, they must be the reason (and the only reason) the agent didn't save the child. I think it's clear that the sharks play no such causal role in the actual sequence of events, given how the case is set up.

Of course, there are other possible scenarios where the sharks (and only the sharks) do cause you to refrain from jumping into the water and saving the child. In those scenarios you omit to act because you couldn't have done otherwise, in Frankfurt's sense. Of course, in those scenarios you are not morally responsible for not saving the child either. But recall that PPA-revised is supposed to capture the *whole* truth about the relationship between moral responsibility for omissions and alternative possibilities. Scenarios like Sharks show that it fails to do so. For you are not morally responsible for your omission in Sharks, your lack of moral responsibility is in some important way connected to your inability to do otherwise, but PPA-revised fails to explain why you are not morally responsible for your omission in that case.

By the way, this is interesting because it suggests that, even if Frankfurt were right that actions and omissions are symmetrical in that being morally responsible for them doesn't require having the ability to do otherwise, i.e. they are symmetrical with respect to what the relationship between responsibility and alternative possibilities *isn't*, an asymmetry would still crop up between them when trying to explain what the relationship between responsibility

and alternative possibilities *is* (given that it's not that responsibility requires the ability to do otherwise). Let's say that Frankfurt is right and, in cases of action, the relationship consists in the fact that responsibility requires not acting only because one couldn't have done otherwise. It would seem that, in cases of omission, the relationship does not consist in this fact, or is not exhausted by this fact. You fail to be morally responsible if the sharks cause you to omit to act, but you *also* fail to be morally responsible if the sharks just happen to be there and don't cause you to omit to act. In other words, the connection between moral responsibility and alternative possibilities is particularly robust in the case of omissions, in a way that escapes Frankfurt's principle.⁷

3. The phenomenon of mutual causal cancellation

So let's ask once again: Why aren't you morally responsible for not saving the child in Sharks? As we have seen, it won't help to say that the inevitability factors made you not save the child: the inevitability factors didn't actually make you do (or not do) anything, because they were causally inoperative. But they still seemed to play *some* role. What role did they play? As we will see, reflecting on the role played by the sharks will help us solve the billionaire puzzle.

Here's an idea about the sharks' role that seems, on the face of it, more promising. Despite their being causally inefficacious (or, maybe: in addition to their being themselves causally inefficacious), the sharks were responsible for the fact that *another* factor was also causally inefficacious. Which other factor? Assuming a broadly causalist framework of agency, one could say: whichever factor is the cause of an agent's not acting when the agent intentionally omits to act. Presumably, this will be some mental event or state concerning the

agent. Following many causalists, let's say that it's the agent's having formed an intention, such as the intention to omit to perform the act in question.⁸ Then the relevant factor in Sharks is your having formed the intention not to save the child. Had it not been for the fact that the sharks were present, you would have been able to carry out your intention not to save the child: your intention not to save the child would have then accounted for, or it would have caused, your not saving the child, and thus you would have been morally responsible for not saving him. But, in fact, given the presence of the sharks, your intention not to save the child didn't account for, or didn't cause, your not saving the child. As a result, you weren't morally responsible for not saving him.

In other words, according to this proposal, the role played by the sharks is that the sharks severed the causal relationship that would otherwise have existed between your intention and your omission, without themselves causing your omission. It will be useful to distinguish *three* different features that the sharks have vis-à-vis the causal history of your omission to save the child. First, an actual (negative) causal feature: the sharks are in fact *not* a cause of your not saving the child. Second, a counterfactual (positive) causal feature: the sharks *would* have been a cause of your not saving the child if you had intended to perform the relevant act (if you had formed a different intention). And third, a feature that determines the (negative) causal role of other factors: the sharks are the reason your intention not to act isn't a cause of your not saving the child *either*.

The role played by the sharks is analogous to the role that certain factors play in a particular kind of causal structure that has been recently discussed in the literature on causation. Here is an example of it:

The Catcher and the Wall: Someone throws a baseball and you catch it. Behind you there is a solid brick wall, which would have stopped the ball if you hadn't caught it. Behind the wall, there is a window, which remains intact throughout that time.⁹

Consider the question: What role does *the wall* play in this case? The first thing to note is that, given that you caught the ball and thus it never touched the wall, it seems wrong to say that the wall is responsible for the window's remaining intact in this case (or that the wall actively prevents the window from breaking). The wall never gets to do anything, in particular, it never gets to deflect the ball away from the window.

But, then, what role does the wall play in this situation? It is plausible to say the following: the wall "protected" the window but only in the sense that, due to its existence, the window was never in danger of breaking. As a result, other things *also* failed to cause the window's remaining intact. In particular, *your catch* failed to cause the window to remain intact. In other words: the wall rendered your catch causally inefficacious, and it rendered your catch causally inefficacious by making the window "unbreakable."

Again, we can distinguish three different features that the wall has concerning the causal history of the outcome of the window remaining intact. First, the wall didn't cause the window to remain intact. Second, the wall would have caused the window to remain intact if you hadn't caught the ball (since, if you hadn't caught the ball, the window's remaining intact would have depended on the wall's presence). And, third, the wall is the reason that your catch didn't cause the window to remain intact either (if the wall hadn't been there, your catch would have been a cause of the window's remaining intact). So the wall in this case is the

analogue of the sharks in *Sharks*; it plays exactly the same kind of role vis-à-vis the causal history of the outcome in question.

The structure of *The Catcher and the Wall* is usually called “preemptive prevention” in the literature on causation. However, this isn’t a good label, at least not in this particular case. “Preemptive prevention” suggests that there is an outcome that is being prevented, a preempting factor that actually does the preventing, and a preempted factor that would have done the preventing in the absence of the preempting factor. *The Catcher and the Wall* would be a case of preemptive prevention in this sense if the catcher were the preempting factor and the wall were the preempted factor with respect to the prevention of the window shattering. But, as we have seen, this is not what happens: the wall isn’t what prevents the window from shattering, and the catcher isn’t what prevents the window from shattering either. “Mutual causal cancellation (neutralization)” or perhaps, more specifically, “mutual prevention cancellation (neutralization),” are more appropriate labels in this case, since they capture the fact that neither is a cause of the relevant non-occurrence (neither is a “preventor”) and that this is so due to each other’s presence (the wall isn’t a preventor given what the catcher does and, vice-versa, the catcher isn’t a preventor given the wall’s existence). So these are the labels I will use.¹⁰

4. Solving the billionaire puzzle

Now let us return to our puzzle about intentional omissions. How can the preceding discussion help us solve the billionaire puzzle?

Basically, the proposal has two parts. The first part is to suggest that the glue tying the chair to the ground in the billionaire case is like the sharks in *Sharks* and like the wall in *The*

Catcher and the Wall. Although the glue didn't itself account for, or didn't cause, your not moving the chair (it would have, if you had tried to move the chair, but it didn't in the actual scenario), it was responsible for the fact that the relevant mental events/states concerning you (e.g., your having formed the intention not to move the chair) weren't causally connected to your not moving the chair. This would be yet another illustration of the phenomenon of mutual causal cancellation or neutralization: the glue didn't cause your not moving the chair because you didn't intend to move it, but, also, the relevant mental items concerning you didn't cause your not moving the chair because of the glue. In other words, the inevitability factors and the relevant mental items concerning the agent render each other causally inefficacious, in that they deny each other the opportunity to be active preventors. Then the second part of the proposal is to suggest that this explains why your not moving the chair wasn't intentional. For, again, following a broadly causalist perspective of agency, one could suggest that an omission is not intentional unless it was (suitably) caused by the relevant mental items concerning the agent. For example, one could suggest that, even if you formed the intention not to move the chair, and even if you didn't move the chair, you couldn't carry out your intention not to move the chair because the intention didn't cause your not moving the chair. And this was due to the presence of the inevitability factors. Hence, to the extent that the inevitability factors are responsible for severing that causal relationship, they are responsible for the fact that the non-doing isn't an intentional omission.¹¹

Now, before we take this idea any further, let me make one important clarification. As it should be clear by now, I think that the reason why you don't intentionally omit to move the chair in the billionaire case is basically the same reason why you are not morally responsible for your not saving the child in Sharks. But I don't thereby mean to suggest that the

relationship between moral responsibility for omissions and alternative possibilities is *exactly the same relationship as* that between intentionally omitting to act and alternative possibilities (in fact, I think it's probably not, as I explain momentarily). All I mean to say is that sometimes the fact that an agent isn't morally responsible for an omission and the fact that an agent doesn't intentionally omit to act can have the same source. Intuitively, this is true of Sharks and the billionaire case: intuitively, the reason you don't intentionally omit to move the chair in the billionaire case is the same kind of reason why you fail to be morally responsible for your omission in Sharks, namely, that you wouldn't have been able to perform the relevant act even if you had tried. There could be other kinds of cases where an agent fails to be morally responsible for an omission in virtue of his inability to do otherwise but where he still intentionally omits to act. This might be true, for example, of a scenario of coercion. Imagine that someone issues a powerful threat against my saving the child (a threat that any reasonable person would succumb to, e.g., he threatens to kill my family if I do), but I can save him if I decide to do so. Imagine that I omit to save the child as a result of the threat (and only as a result of the threat). In this case, I intentionally omit to save the child, but (at least Frankfurt would say) I am not responsible for omitting to save him. Moreover, my lack of moral responsibility in this case too seems to be importantly tied to my inability to do otherwise. Scenarios of this kind suggest that the relationship between moral responsibility for omissions and alternative possibilities is probably not the same as the relationship between intentionally omitting to act and alternative possibilities. But this isn't a problem for our argument here.¹²

Now we are ready to state the solution to the billionaire puzzle in more precise terms. Part of the solution consists in the formulation of a new principle about the connection between intentionally omitting to act and alternative possibilities. This principle is:

PMCC-IO (“Principle of Mutual Causal Cancellation for Intentional Omissions”): An agent does not intentionally omit to perform an act (even if he doesn’t perform the act, and even if he has the relevant desires, intentions, etc. not to act) if inevitability factors and the relevant mental items enter in a relation of mutual causal cancellation with respect to the agent’s not acting.

The other part of the solution is the claim that the billionaire case is a scenario of mutual causal cancellation of the relevant kind. The key to solving the puzzle, then, is to switch the focus from what *moves* the agent to omit to act to what *doesn’t* move the agent to omit to act. As suggested by PMCC-IO, inevitability factors can play a role in rendering the omission non-intentional even if they don’t themselves move the agent to omit to act: in particular, they are responsible for the omission’s not being intentional when it’s in virtue of *them* that the relevant mental items concerning the agent don’t result in the agent’s omission.¹³

I will end this section by making two remarks about PMCC-IO: one of them is friendly to Frankfurt’s ideas and the other one isn’t. First, the Frankfurt-friendly remark: It is important to see that PMCC-IO is consistent with the claim that an agent can intentionally omit to act when he couldn’t have done otherwise (in other words, PMCC-IO is not PPA-IO). According to PMCC-IO, agents cannot intentionally omit to do the impossible when the inevitability factors and the relevant mental items render each other causally inefficacious.

But, of course, for all PMCC-IO says, agents *can* intentionally omit to do the impossible to the extent that this isn't the case. In particular, PMCC-IO is consistent with the claim that, in a Frankfurt-style omission case, the agent intentionally omits to act even if he couldn't have done otherwise. Note that the structure of a Frankfurt-style omission case is such that the agent freely comes to form an intention and then omits to act, based on that intention. The counterfactual intervener is situated, so to speak, "before the intention" rather than "after the intention": he would have prevented the agent from forming a different intention, rather than preventing him from doing anything different *once* he had formed a different intention. So in a Frankfurt-style omission case there is no temptation to say that the inevitability factors sever the link between the agent's intention and his omission (and thus there is no temptation to say that the intention and the inevitability factors render each other causally inefficacious). Although the agent couldn't have formed a different intention, arguably, (if intentions to omit to act ever result in agents' omissions) the intention that the agent formed did result in his omission. For example, an advocate of Frankfurt-style cases would claim that, in The Child and the Neuroscientist case (our Frankfurt-style omission scenario from section 1), you freely form the intention not to jump in to save the drowning child, and the child drowns as a result of your forming that intention. Although the existence of the neuroscientist makes it impossible for you to decide to save the child, it doesn't cut off whatever connection exists between your intention not to save the child and your not saving him (and nothing else does, in particular, there are no sharks).¹⁴

Here is another way to see this: an advocate of Frankfurt-style cases could argue that The Child and the Neuroscientist doesn't have the structure of a mutual causal cancellation

scenario, like *The Catcher and The Wall*. Instead, an advocate of Frankfurt-style cases could argue that it has the structure of a scenario like the following:

The Catcher and the Neuroscientist: Someone throws a baseball and you catch it. Behind you there is a window, which would have shattered if the ball had reached it. This time there is no brick wall or any other obstacles to the ball's reaching the window. But, as it turns out, you couldn't have failed to catch the ball. A neuroscientist has been closely monitoring your movements: if you hadn't instinctively placed your glove in the right place a few seconds before the ball got there, he would have sent some signals to your brain that would have resulted in your catching the ball all the same.

Here we don't hesitate to think that your catch caused the window to remain intact. Of course, you couldn't have failed to catch the ball, but this is no objection to the claim that your catch prevented the shattering. Similarly, in *The Child and the Neuroscientist* you couldn't have formed the intention to jump in, but, it could be argued, this is no objection to the claim that your intention not to jump in caused your omission to save the child.

Now for the remark about PMCC-IO that is not Frankfurt-friendly: PMCC-IO is consistent with there being an asymmetry between actions and omissions concerning the conditions for acting/omitting to act intentionally.¹⁵ For all PMCC-IO says, inevitability factors might sever the link between the relevant mental items and the agent's behavior in cases of omission but not in cases of action. Or they might sever that link much more often in cases of omission than in cases of action. That is, it might be that, although both acting

intentionally and omitting intentionally to act require that the relevant mental items cause the behavior, the conditions under which the relevant mental items cause the behavior are different for actions and omissions, given the different role played by inevitability factors in each case. Elsewhere I have argued that there is, in fact, a causal asymmetry between actions and omissions concerning the role of inevitability factors.¹⁶ If so, this fact, together with PMCC-IO, would entail that the relationship between intentional omissions and alternative possibilities is ultimately different from the relationship between actions and alternative possibilities.

5. Final Remarks

According to PMCC-IO, *if* and *when* there is mutual causal cancellation between inevitability factors and the relevant mental items concerning the agent, the agent doesn't intentionally omit to act. Although PMCC-IO doesn't specify any particular conditions under which there is mutual causal cancellation between the inevitability factors and the relevant mental items, I have argued that there are at least some paradigmatic cases where this happens. In particular, I have argued that this is what happens in the billionaire case and in other scenarios with the same structure. So PMCC-IO, in conjunction with the claim that the billionaire case is a scenario of mutual causal cancellation, provides a solution to the billionaire puzzle.

I said at the outset that reflecting about the billionaire puzzle would be enlightening even if it turned out that intentional omissions do require alternative possibilities (in particular, even if Frankfurt-style cases failed to show that they don't). Now we can see why. Reflecting about the billionaire puzzle led us to the formulation of PMCC-IO, an alternative principle about the relationship between intentional omissions and alternative possibilities

(one that doesn't presuppose that intentional omissions require alternative possibilities). As I noted at the outset, even if intentional omissions did require alternative possibilities, that is, even if PPA-IO were, after all, true, there would still be a question as to why it is true. And it seems that, if PPA-IO were true, we could appeal to PMCC-IO to explain why it is true. For, given PMCC-IO, if PPA-IO were true, this would have to be because in *all* cases of omission where inevitability factors are present, the inevitability factors and the relevant mental items enter in a relation of mutual causal cancellation. If this were true, it would explain why intentional omissions require alternative possibilities: they would require alternative possibilities because they would require the absence of inevitability factors, and they would require the absence of inevitability factors because, were there to exist any inevitability factors, they would (by means of entering in a relation of mutual causal cancellation with the relevant mental items) break off the causal links that are essential for an intentional omission to obtain).¹⁷

Let me conclude by describing one important way in which I think that the project in this paper might be incomplete, or might need some more pruning. I take it that the billionaire case (and other scenarios with a similar structure) are paradigm examples of omission where the agent's inability to act precludes his having intentionally failed to act. I've offered an account of this in terms of the phenomenon of mutual causal cancellation. Now, on the face of it, there could be other examples that fall out of this account because they don't involve causal connections (or the lack thereof). Imagine, for instance, an arrogant mathematician who believes that a certain mathematical claim is a theorem but doesn't bother to try to prove it. Imagine that the claim isn't really a theorem (it cannot be proved from the axioms). Does the mathematician intentionally omit to prove it? Presumably not, since he couldn't have proved

it. So the mathematician's not proving the theorem fails to be an intentional omission, just like your not moving the chair in the billionaire case does. Could we tell the same kind of story I've told about the billionaire case in the mathematician case? Arguably, it couldn't be exactly the same story. One feature of the billionaire case, in virtue of which I have claimed there is mutual causal cancellation in that case, is that the inevitability factor in question (the glue tying the chair to the ground) would have been causally responsible for your not moving the chair if you had tried to move it. But in the mathematician case the inevitability factor is a mathematical fact: the fact that a particular mathematical claim doesn't follow from certain axioms. Presumably, a mathematical fact couldn't have been *causally* responsible for the mathematician's failure to prove the theorem, if he had tried to prove it. If so, there isn't mutual causal cancellation in this case.

I see two main ways to go about addressing this problem. First, one could argue that what's at issue here is a broader kind of mutual cancellation: mutual *explanatory* cancellation, which needn't be of the causal kind. On this view, had the mathematician tried to prove the claim, the fact that the claim isn't a theorem would still (non-causally) explain why he fails to do so. (Note that this is what one would be tempted to say anyway if one believed, contrary to what I have assumed in this paper, that omissions cannot be caused by anything or be causes of anything; in that case, omissions couldn't give rise to situations of mutual causal cancellation, but they would still likely give rise to situations of mutual explanatory cancellation.) Second, one could argue that (at least sometimes) agents cannot genuinely try to bring about states of affairs that are logically impossible, e.g., I cannot try to bring it about that $p \ \& \ \text{not-}p$.¹⁸ If so, perhaps the fact that I cannot even try to bring it about that $p \ \& \ \text{not-}p$ more directly explains why I cannot intentionally fail to bring it about that $p \ \& \ \text{not-}p$. Either

way, the story would have to be modified or complemented accordingly, if we wanted to account for this kind of cases. At any rate, I suspect that cases of this kind are somewhat special and out of the ordinary, and thus they probably deserve a special treatment, which I cannot give them here.

References

- Albritton, R. (1985) "Freedom of Will and Freedom of Action," Presidential Address, American Philosophical Association (Pacific Division), reprinted in G. Watson (ed.) *Free Will*, Oxford U. P., second edition, 2003, pp. 408-23.
- Bach, K. (2010) "Refraining, Omitting, and Negative Acts," in T. O'Connor and C. Sandis (eds.), *Companion to the Philosophy of Action*, Oxford: Wiley-Blackwell.
- Clarke, R. (1994) "Ability and Responsibility for Omissions," *Philosophical Studies* 73: 195-208.
- Collins, J. (2000) "Preemptive Prevention," *Journal of Philosophy* 97, 4: 223-34.
- Fischer, J. and M. Ravizza (1991) "Responsibility and Inevitability," *Ethics* 101: 258-78.
- Fischer, J. and M. Ravizza (1998) *Responsibility and Control*, Cambridge U. P.
- Frankfurt, H. (1969) "Alternate Possibilities and Moral Responsibility," *Journal of Philosophy* 66, 23: 829-39.
- Frankfurt, H. (1994) "An Alleged Asymmetry Between Actions and Omissions," *Ethics* 104, 3: 620-3.
- Frankfurt, H. (2006) "Some Thoughts Concerning PAP," in D. Widerker and M. McKenna (eds.), *Moral Responsibility and Alternative Possibilities*, Ashgate, pp. 339-45.

- Ginet, C. (2004) "Intentionally Doing and Intentionally Not Doing," *Philosophical Topics* 32, 1-2: 95-110.
- Ginet, C. (2006) "In Defense of the Principle of Alternative Possibilities: Why I Don't Find Frankfurt's Argument Convincing," in D. Widerker and M. McKenna (eds.), *Moral Responsibility and Alternative Possibilities*, Ashgate, pp. 75-90.
- Locke, J. (1975) *An Essay Concerning Human Understanding*, ed. P. Niddich, Oxford U. P.
- McDermott, M. (1995) "Redundant Causation," *British Journal for the Philosophy of Science* 46: 523-44.
- McIntyre, A. (1994) "Compatibilists Could Have Done Otherwise: Responsibility and Negative Agency," *Philosophical Review* 103: 453-88.
- Sartorio, C. (2005) "A New Asymmetry between Actions and Omissions," *Noûs* 39, 3: 460-82.
- Sartorio, C. (2006) "Disjunctive Causes," *Journal of Philosophy* 103, 10: 521-38.
- Sartorio, C. (2009) "Omissions and Causalism," *Noûs* 43, 3: 513-30.
- van Inwagen, P. (1978) "Ability and Responsibility," *Philosophical Review* 87: 201-24.
- Widerker, D. (2000) "Frankfurt's Attack on the Principle of Alternative Possibilities: A Further Look," *Philosophical Perspectives* 14: 181-201.
- Zimmermann, M. (1981) "Taking Some of the Mystery out of Omissions," *Southern Journal of Philosophy* 19: 541-54.

* I am grateful to Kent Bach, Juan Comesaña, Dan Hausman, Ori Simchen, and audiences at the University of North Carolina-Chapel Hill, the University of British Columbia, and Torcuato Di Tella University for helpful comments on earlier versions of this paper.

¹ The reason I say that it might seem even clearer that you intentionally omit to save the child in this case (than that you are morally responsible for your omission) is that, one could argue, you would have intentionally omitted to save the child *even if* the neuroscientist had intervened. For, given the way the case is set up, even if the neuroscientist intervenes, he intervenes by causing you to decide not to save the child; as a result, you still fail to save the child based on your decision not to save him. For our purposes, we don't need to settle the

question whether your not saving the child would be intentional in that case too; all that is needed is the claim that your omission is intentional in the scenario where the neuroscientist doesn't intervene.

² See Sartorio (2005) (and see Sartorio (2009) for complementary arguments).

³ See also, e.g., Clarke (1994) and McIntyre (1994).

⁴ Frankfurt writes: "In my view, there is every reason to prefer an account that is straightforwardly symmetrical." And: "The two cases [the relevant omission and action cases under his consideration] are, in any event, perfectly symmetrical. If there is any discordance in holding John [the agent in a given omission case] morally responsible for refraining from saving the child, it is equally discordant to regard Matthew [the agent in the parallel action case] as morally responsible for saving him." (Frankfurt (1994), pp. 622-3.)

⁵ He writes: "Now it seems to me that when we turn from cases of action to cases of omission, we find the very same possibility. In the actual sequence of events, the fact that someone could not have performed a certain action—and hence could not have avoided omitting it—may have no causal influence on his behavior. It may play no role whatsoever in accounting for his omitting to perform the action." (Frankfurt (1994), p. 621) Here Frankfurt is thinking of cases where the agent is morally responsible for an omission even if he couldn't have performed the relevant act. He seems to think that part of the explanation of why an agent is responsible in these cases is that he doesn't omit to act *because* he couldn't have done otherwise.

⁶ Ginet makes a similar point about a different example in his (2006), pp. 82-3. See also Widerker (2000), pp. 189-90.

⁷ Compare with a parallel case of action: The child is placidly napping on a raft in the water. You want the child to drown so you create waves by violently moving your body in the water. The child falls off the raft and drowns. Had you not independently decided to create the waves, you would have noticed the sharks in the water, which would have resulted in your swimming quickly towards the shore, which would have, again, resulted in the waves and the child drowning. In this case the mere presence of the sharks doesn't get you off the hook: you would have failed to be morally responsible for making the child drown if the sharks had caused you to create the waves, but you are morally responsible given that they didn't.

A different possible response by Frankfurt is suggested in his (2006), pp. 341-2. The response consists in saying that the agent in Sharks is not morally responsible for the omission to save the child simply because there is no such omission: there is no such omission on the part of the agent since saving the child wasn't in his control. I think that this response would only push the problem back one step. For then the question would be: Why is there no omission in Sharks for the agent to be morally responsible for, but there is one in a Frankfurt-style case, if in both cases the agent couldn't have performed the relevant act?

⁸ Contrary to what many causalists seem to believe, I don't think that intentions not to act are the relevant mental causes in cases of intentional omission (see Sartorio (2009)); the view that I favor thus ends up being an unorthodox version of causalism, and only if one understands "causalism" very broadly, as I am understanding it here). But I still think that the sharks render the relevant mental items causally inefficacious in this case. So what exactly one takes the relevant mental items to be in this case is not essential for the point about Sharks to hold.

⁹ The example is originally by McDermott (1995)). I am interested in the treatment given of it by Collins (2000), which differs to an important extent from McDermott's.

¹⁰ Thanks to Kent Bach for an illuminating discussion of the appropriate labels to use in these cases. The reason this type of structure is usually called "preemptive prevention" is that there are other examples with apparently the same basic structure where it seems more plausible to say that one of the factors preempts the other factor by actually doing the preventing (for example, if we replace the solid brick wall with another human catcher; in that case it seems more plausible to say that the agent who catches the ball prevents the shattering—see Collins (2000) for the contrast between the two cases). Unlike genuine scenarios of preemptive prevention, mutual causal cancellation scenarios seem to give rise to a puzzle. In the Catcher and the Wall, had both the catcher and the wall been absent, the window would have shattered. So it seems that the catcher and the wall "together" prevented the shattering. But how can this be, if neither was a preventor? A possible answer is to say that a disjunctive fact (the fact that either was present) did the preventing. Notice that a similar puzzle would arise for Sharks: if the sharks didn't cause the child's death, and neither did I, then what did? For more reasons to believe in disjunctive causes, see Sartorio (2006).

¹¹ The same can be said about Locke's famous example of the man who stays inside a room talking to someone he likes, although, unbeknownst to him, the door has been locked from the outside so that he cannot get out (Locke (1975), bk. 2, ch. 21, sect. 10). It seems to me that, although the man intentionally omits to *try* to leave

the room, he doesn't intentionally omit to leave the room, for the same reason you don't intentionally omit to move the chair in the billionaire case.

¹² Another kind of scenario where the agent is not morally responsible for his omission because he couldn't have done otherwise but where, it could be argued, he still intentionally omits to act is a Frankfurt-style scenario where the neuroscientist intervenes by causing the agent to decide not to act in the relevant way (see n. 1 above). In light of cases like these, a tentative way to formulate the principle about the connection between moral responsibility for omissions and alternative possibilities could be: one is not morally responsible for an omission *either* if one's intention not to act was caused by inevitability factors *or* if inevitability factors and one's intention rendered each other causally inefficacious with respect to the fact that one didn't act.

¹³ A similar shift in focus would help explain why the agent isn't morally responsible in Sharks. As I pointed out in section 2, Frankfurt's guiding thought about the connection between moral responsibility and alternative possibilities was that inevitability factors can only reduce the agent's moral responsibility for his act/omission if they play an actual role in making the agent act/omit to act. In light of our preceding discussion, I think we should say that inevitability factors can also reduce the agent's responsibility if they cut off the connection between the relevant mental items concerning the agent and his act/omission *without themselves making him act/omit to act*. This is so because, in that kind of case too, the agent doesn't act/omit to act *for his own reasons*.

¹⁴ Clarke (1994) and McIntyre (1994) suggest something along these lines as part of their justification for the claim that the agent is morally responsible for his omission in cases with this structure. Note that the claim that the agent's intention results in his omission in these cases could still fail if it were *generally* false that the intention results in the omission in cases of intentional omission. But what we cannot say, it seems to me, is that *the inevitability factors* are the reason the intention doesn't result in the omission in these cases.

¹⁵ This shouldn't be surprising, given the divergence between PMCC-IO and Frankfurt's own principle (PPA-revised), discussed in section 2.

¹⁶ See Sartorio (2005). There I draw attention to a moral asymmetry that results between actions and omissions as a result of the causal asymmetry, not to an asymmetry concerning the conditions for acting/omitting to act intentionally.

¹⁷ Notice that the claim would have to be that Frankfurt-style cases are scenarios of mutual causal cancellation too, just like Sharks or the billionaire case. But how could this be, if, as we have seen, it is not plausible to say that in Frankfurt-style cases the inevitability factors sever the link between the agent's intention not to act and his failing to act? As suggested above (see notes 8 and 14), one would have to say that the agent's intention not to act is not the relevant mental item, in this or other omission cases. I have argued for this in Sartorio (2009). As I explain there, I take this to be a rather surprising fact about omissions, and one that is definitely not immediately obvious. The fact that it's not immediately obvious could be used to "explain away" the initial appearance that an agent intentionally omits to act in a Frankfurt-style case.

¹⁸ Albritton defends a view of this kind in his (1985).