

# Failures to Act and Failures of Additivity\*

Carolina Sartorio

University of Wisconsin-Madison

## 1. Introduction

On the face of it, causal responsibility seems to be “additive” in the following sense: if I cause some effects, then it seems that I also cause the sum (aggregate, conjunction, etc.) of those effects. Let’s call the claim that causation behaves in this way, Additivity.<sup>1</sup>

The intuitive appeal of Additivity can be illustrated by the following example. Imagine that I launch an attack on someone by simultaneously kicking him and punching him, and that my kicks cause certain bruises on my victim’s body (say, on his legs) whereas my punches cause certain other bruises (say, on his face). Then it seems that I cause the sum of those bruises (the face and leg bruises). Intuitively, this generalizes: if something I do issues in part of an outcome and another thing I do issues in another part of the outcome, then it seems that we should say that the sum of my actions issues in the outcome as a whole and thus, as a result, that I am causally responsible for the outcome by virtue of having acted in those ways.<sup>2</sup>

Now, it also seems that agents can cause outcomes by omission, i.e. not by acting in certain ways but by failing to act in certain ways.<sup>3</sup> Hence, one would expect Additivity to apply to omissions as well. In the case of omissions, Additivity would claim that, if I cause some effects by failing to act in certain ways, then I also cause the sum (aggregate, conjunction, etc.) of those effects. Again, this seems intuitively right. The motivating thought is the same as in the case of (positive) actions: if something I fail to do is causally responsible for a part of an outcome and another thing I fail to do is causally responsible for another part of that outcome,

then it seems that the sum of my omissions must be responsible for the outcome as a whole, and thus I must be responsible for the outcome in virtue of having failed to act in those ways. If we conceive of omissions as facts, then the sum of two omissions will be a conjunctive fact. For instance, the sum of my failing to talk and my failing to walk is the fact that I didn't talk and I didn't walk. Hence the Additivity principle for omissions is the claim that, if each of a number of omissions of mine makes me causally responsible for an outcome, then the fact that I failed to act in each of those ways makes me causally responsible for the sum of those outcomes.<sup>4</sup>

I will argue, however, that Additivity fails for omissions. This paper has three parts. First, I offer counterexamples to Additivity involving omissions. Second, I consider and rebut some objections to the alleged counterexamples. Third, I argue that it is actually a good thing that Additivity fails in those cases, and thus we should be relieved that it does.

## **2. Counterexamples to Additivity**

I will offer two counterexamples to Additivity involving omissions.

Case 1:

Battlefield: I am at the battlefield and I see that some of our soldiers are about to be slaughtered by the enemy. I could save any one of them, but only one of them (I only have one bullet left). I cannot get myself to choose which one to save so they all die.

Consider the claim "I caused all of those soldier's deaths (by failing to act in the relevant ways)." Is this true or false? Our first reaction is to say that it's false: I could have saved at most one soldier, so I am not a cause of all their deaths.<sup>5</sup> So far so good. But now consider the following

reasoning. Take any soldier in that group. It is true, of that soldier, that I could have saved him. Hence, given that I didn't save him, I am among the causes of his death. This is true of all the soldiers in that group. Therefore, doesn't it follow, after all, that I caused all of those soldier's deaths?<sup>6</sup>

In other words, Battlefield seems to give rise to a puzzle: on the one hand, there is reason to say that I didn't cause all of the soldiers' deaths but, on the other hand, there is reason to say that I did. A natural way out of the puzzle is to say that there is a reading of the statement "I caused all of those deaths" on which it is true and a reading on which it isn't. The reading on which it is true is the distributive reading: "I caused all of those deaths" means "I caused each of those deaths." And the reading on which it is false is the collective reading: "I caused all of those deaths" means "I caused the sum of those deaths."<sup>7</sup> In other words, Battlefield is a counterexample to Additivity: although I caused each of the deaths, I didn't cause their sum.<sup>8</sup>

Case 2:

Fire: A fire breaks out in the chemical plant where I work. The building consists of several rooms. I only have time to seal (and thus to save) one of the rooms. Instead, I run out, and the whole building is destroyed.

Consider the claim "I caused the whole building to be destroyed." Again, this strikes us as false. But, again, there is a reading of the claim on which it seems that we have to say that it is true (the distributive reading): for each of the building's rooms, I could have saved it and I didn't, so I was a cause of its destruction. Hence, for each of the rooms, I caused its destruction, but I didn't cause the sum of those rooms being destroyed. Again, Additivity fails.

At this point, one might ask: Isn't it possible to build a similar challenge to Additivity involving commissions? In particular, isn't it possible to devise a version of the kicking-and-punching example from last section where Additivity fails, and for the same reasons? The new version would have to look as follows. I simultaneously kick and punch my victim, thus causing bruises on his legs and face; I could have refrained from kicking him and I could have refrained from punching him; however, I couldn't have both refrained from kicking him and punching him (say, an evil neuroscientist who is controlling my behavior would have made sure that I acted in at least one of those ways). In this case it seems that I cause the bruises on my victim's legs and I cause the bruises on my victim's face, but I also cause the sum of those bruises. The claim "I caused all of those bruises" doesn't seem to have a reading on which it is false. Hence, this case doesn't violate Additivity. Now, one might wonder whether this generalizes to all cases involving commissions. I will not attempt to answer this question here; this at least shows that it is much harder to build a case against Additivity with commissions than with omissions. Thus I will rest my case on counterexamples involving omissions, like Battlefield and Fire.<sup>9</sup>

In the next section I anticipate some objections to these alleged counterexamples and I try to rebut them.

### **3. Objections and Replies**

#### **Objection 1: Sophie's choice**

Recall Sophie: she is given the difficult choice of deciding which, if any, of her two children will survive (otherwise they'll both die). Imagine that Sophie couldn't get herself to choose and both died. Wouldn't we say then that she caused both of her children to die? If so, wouldn't we have to say that she caused the sum of her two children's deaths? But this case has

the same structure as the examples I have offered, so it seems that if Sophie causes the sum, I cause the sum in my examples too.

Reply: Again, the claim “Sophie caused both of her children to die” is ambiguous between “She caused each of their deaths” and “She caused the sum of their deaths.” My claim is that she does cause both of their deaths in the first sense, but not in the second sense. Imagine that we said, disapprovingly, of Sophie: “She didn’t do the right thing. She caused both of her children to die!”. I don’t think we would do that to emphasize that we are referring to the sum of their deaths, but only to point out that she could have saved one of her children and didn’t.

#### Objection 2: Causing the part is causing the whole

Here is a plausible principle: causing a part is sufficient for causing the whole, in the relevant sense of “cause” which is contributory cause (as opposed to the cause, or the only cause, or a sufficient cause). The underlying thought behind this principle is: suppose you cause a part, then you’re a step (even if it’s a tiny step) closer to causing the whole. Thus whatever causes a part contributes towards, and thus is a cause of, the whole. Now, take, e.g., the Battlefield case. Consider any soldier whose life I could have saved at the expense of not saving the life of others. I grant that I caused that soldier’s death (that is, I contributed to that death). But then, given that the death of that soldier is part of the sum of all the deaths, the principle that causing the part is sufficient for causing the whole entails that I also caused the sum (that is, I contributed to the sum).

Reply: This reasoning fails, for the following reason. By assumption, in this case I caused all of the parts, not just one: I contributed to each of the soldiers’ deaths. So, according to the above reasoning, my contribution shouldn’t count as one tiny step towards the whole but as

many such steps. But then, if the reasoning were sound, it would follow that I am much more of a contributor than I can plausibly be said to be. To see this, compare my contribution to that of someone whose job was to load each of the soldiers' guns but failed to do so (as a result, they couldn't defend themselves). He is someone who contributed to the sum of the soldiers' deaths in a way that can be regarded as many tiny steps towards that outcome. My contributions to the individual deaths, by contrast, don't seem to add up in this way. Failing to save one soldier's life doesn't add up to failing to save another soldier's life when I couldn't have saved more than one at a time.

Here is another way to bring out the disanalogy. Given that I couldn't have saved more than one life, I am not a little bit worse by having failed to save two lives instead of just one. By contrast, whoever failed to load the guns is a little bit worse by having failed to load two of the guns instead of just one. This suggests that the tiny steps reasoning only works when and if causation is additive, otherwise it doesn't.

### Objection 3: When there is counterfactual dependence there is causation

Take the Battlefield case again. By assumption, I could have saved one of the soldiers (any one). Had I saved one of the soldiers (any one), the actual sum of deaths wouldn't have occurred; instead, a different sum would have occurred (that consisting of the deaths of those soldiers who die in the counterfactual scenario). In other words, the sum that actually occurs counterfactually depended on each of my failures to save one of the soldiers. But, presumably, counterfactual dependence is sufficient for causation. If an outcome would not have occurred had it not been for something I did or failed to do, then it seems that the outcome occurred because of what I did or failed to do, and thus, presumably, what I did or failed to do caused the outcome.

For instance, imagine that, had I not been startled by a noise outside, I wouldn't have spilled my coffee. In that case we would think that my being startled by the noise caused my spilling my coffee. Or imagine that, had I called Grandma on her birthday, she wouldn't have felt sad. In that case we would think that my failure to call her caused her to be sad.<sup>10</sup> Now, if counterfactual dependence is sufficient for causation, then it follows that, in Battlefield, I was a cause of the sum of the soldiers' deaths.

Reply: The claim that counterfactual dependence is sufficient for causation is controversial. In the literature on causation, several objections have been pressed against that claim. As an example, it has been pointed out that certain counterfactual connections track non-causal relations such as logical or mereological relations (for instance, my writing "cat" counterfactually depends on my writing "c" but my writing "c" didn't cause my writing "cat"; see Kim (1993)). In addition, some have pointed out that counterfactual connections are too weak to always track genuine causal relations (for instance, Socrates' birth didn't cause his death even if his death counterfactually depended on his birth, or my building a bridge might not cause Jim to cross it even if his crossing it counterfactually depends on my building it; see Thomson (2003)).<sup>11</sup>

Now, the objector might reply that, even if counterfactual dependence is not generally sufficient for causation, it is still evidence of the existence of a causal relation. After all, we do normally rely on the existence of counterfactual connections to settle the status of causal claims. In particular, it seems that the reason we want to say that I caused the individual soldiers' deaths in Battlefield is that it is true of each of those soldiers that, had I acted in the relevant way, he wouldn't have died.

This is fair enough. But I think that we should be wary of appealing to the link between counterfactual dependence and causation to establish the causes of the sum in the cases at hand. For these cases are special, in at least two ways. Let me explain.

First, they share the following feature: for each of the parts of the sum, I caused that part because I could have prevented it and didn't; however, I could only have prevented one part at the expense of not preventing other parts (e.g., the prevention of P1 would have precluded the prevention of P2 and the prevention of P2 would have precluded the prevention of P1). For instance, in Battlefield, it is true of each of the soldiers that I caused that soldier's death because I could have prevented that death and I didn't prevent it; however, I could have only prevented, say, Jim's death at the expense of not preventing Bob's death, and vice versa.

Now, why think that this feature might compromise the link between counterfactual dependence and causation? Here is one possible reason to think so. The motivation for thinking that counterfactual dependence is an indicator of a causal relation is that counterfactual dependence is tied to the concept of "making a difference," and that whenever there is such a thing as difference-making, there is a reason to think there is a causal relation. For instance, the suggestion is that, in Battlefield, I am a cause of the sum of the deaths because I could have made a difference to the sum and I didn't. In particular, I could have helped one of the soldiers (say, Jim), in which case I would have made a difference to the sum by making a difference to one of its parts, Jim's death (because, had Jim's death not occurred, the sum wouldn't have occurred). Similarly for another soldier, say, Bob: I could have helped Bob, in which case I would have made a difference to the sum by making a difference to another part of the sum, Bob's death (because, had Bob's death not occurred, the sum wouldn't have occurred). However, the concept of making a difference is a loose concept. Given the circumstances of the case, had I

helped Jim, I wouldn't have helped Bob. So, even if there is a sense in which, had I helped Jim, I would have made a difference, I would have made a difference in this sense by closing off other opportunities to make a difference (like helping Bob instead). In other words, I could only have made a difference in one sense by failing to make a difference in another sense. Hence, if the main motivation to think that there is a link between counterfactual dependence and causation is the concept of difference-making, there is at least some reason to believe that such a link might fail in cases like Battlefield.

A second reason to believe that the link between counterfactual dependence and causation might break down in these cases is that there are independent reasons to believe that it breaks down in (at least some) cases involving sums. Here is why. Start by considering a simple case of causation, say, an assassin shoots at a victim and the victim dies as a result; moreover, there are no backup assassins or conspiracies. If there is an intimate link between counterfactual dependence and causation, then, presumably, we will want to say that this is a case that illustrates it. That is, we will want to say that the assassin's shooting caused the victim's death, and that this is so (at least partly) because, had the event consisting in the assassin's shooting not occurred, the event consisting in the victim's death would not have occurred either. Now, in order for this to be true, it seems that counterfactual statements must be understood in a specific way (that is, the relation of similarity among possible worlds has to be understood in a specific way). In particular, we should imagine a counterfactual scenario where the actual event of the shooting is not replaced with a barely different event, or by an event that has some of the same bits and pieces of the actual event (e.g., a shooting with a different gun), but one where the actual event is "cleanly and completely excised" from the history of the world (otherwise, if it were

replaced with a barely different event, the death might still have taken place and thus there would be no counterfactual dependence).<sup>12</sup>

Now, how does this apply to sums? Consider the claim: “The sum of the assassin’s shooting and the Pope’s saying a prayer caused the victim’s death.” This seems false.<sup>13</sup> However, based on the above constraint on how to evaluate counterfactuals, the death appears to counterfactually depend on the sum. For, when we take away the sum, we “completely and cleanly” excise it: we also take away its parts. As a result, had the sum of the shooting and the prayer not occurred, then the shooting wouldn’t have occurred and thus the death wouldn’t have occurred either. In other words, the death counterfactually depends on the sum but the sum isn’t a cause of the death. If so, the link between counterfactual dependence and causation breaks down in (at least some) cases involving sums.<sup>14</sup> Hence, we shouldn’t be surprised that it breaks down in cases like Battlefield.

I conclude that the objection from counterfactual dependence fails. The lesson we should learn from it is not that Additivity holds despite cases like Battlefield and Fire but, rather, that such cases constitute (yet) another challenge to the idea that counterfactual dependence is sufficient for causation.

#### Objection 4: Moral versus causal responsibility

The objector might argue that I am confusing moral responsibility with causal responsibility. He might argue as follows. Clearly, in Battlefield, I am not morally responsible for the sum of the soldier’s deaths: given that I couldn’t have prevented more than one death, it seems that I cannot be to blame for their sum.<sup>15</sup> However, I am still causally responsible for it. The reason it strikes us otherwise is that we tend to confuse attributions of moral responsibility

with attributions of causal responsibility (as when we say, of the gardener who failed to water the plant: “He caused the plant’s death, nature didn’t,” although, as a matter of fact, some natural processes also contributed to the death).

Reply: I don’t think our intuitions about these cases are confounded by our moral judgments. To see this, imagine a version of Battlefield where the moral agent is replaced by an automated mechanism. Imagine, for instance, the following variant of the case:

Shield: A defense mechanism is in place in the situation described in Battlefield. At a given time, the mechanism can yield different outcomes (what outcome it actually yields is determined by means of a random process). The different outcomes are: for each soldier, the production of a bulletproof shield that protects that (and only that) soldier, and a null outcome (no shield is created). The mechanism only has the resources to create one shield. As it happens, the null outcome comes about and all of the soldiers die.

Our intuitions about Battlefield are preserved, mutatis mutandis, in Shield. Consider the claim: “The mechanism was causally responsible for all of the deaths.” Again, our first reaction is to say that this claim is false; after all, the mechanism could have saved at most one soldier. But, again, the claim is true on a distributive reading: for each of the soldiers, the mechanism could have created a shield for that soldier and, in that case, that soldier’s life would have been saved. Hence, again, the natural thing to say is that the mechanism is a cause of each of the deaths (individually), but not of their sum. This suggests that our intuitions about, in particular, the original version of Battlefield are not confounded by our moral judgments, but are purely causal intuitions.

As far as I can see, these are the main objections that could be raised against my counterexamples and in defense of Additivity. Since these objections fail, I conclude that we should reject Additivity. In the next section I discuss a consequence of the rejection of Additivity.

#### **4. Failures of Additivity and Famine**

It is a good thing that Additivity fails when it does. In this section I explain why.

As it has been pointed out, the assumption that causation by omission is possible commits us to the belief that there is a lot of it, much more than we would have initially thought.<sup>16</sup> If we want to say, for instance, that the gardener caused the plants' death by failing to water them, it seems that we have to say that everyone else caused their death by failing to water them, including, e.g., the Queen of England. After all, had the Queen of England watered the plants, they wouldn't have died. Similarly, the assumption that causation by omission is possible seems to commit us to the belief that we are causally responsible for the deaths of many starving people in remote places. The thought is: at this very moment, you could be saving lives by writing a check to Oxfam and putting it in the mail, so, if you don't do it, you're causally responsible for those deaths. The fact that those victims are far away, or that they are strangers, or that they are citizens of another country cannot plausibly matter to whether you are a cause of their deaths. Some philosophers go as far as to argue that we are morally responsible for those deaths.<sup>17</sup> There are different ways of trying to resist this serious moral implication. What's important for our purposes here is that, if the causal claim is right (that is, if we really do cause all those deaths), then this is enough to block an otherwise plausible-looking explanation of why we don't take ourselves to be morally responsible for those deaths. (The plausible-looking explanation is: We

don't cause those deaths; so, given that moral responsibility requires causal responsibility, it follows that we are off the hook morally).<sup>18</sup> But, if we do cause those deaths, then we need an alternative explanation of the alleged fact that we are not morally responsible.

Now, if Additivity were true, then our causal involvement vis-à-vis the deaths of people in remote places would be even more unexpected and unwelcome than this. To illustrate this, imagine that I have a certain number of assets and that, in the past year, I could have used those assets in different ways to save the lives of people in remote places. In particular, imagine that there are a million lives such that, for each of them, I could have saved that life by investing my assets in a certain way. However, imagine (as it is reasonable to suppose) that my resources are very limited so I couldn't have saved all of those lives. In particular, imagine that some of those lives are "expensive" lives in that I could only have saved them at the expense of not saving others (say, because they live in very isolated places or they are in a serious state of malnourishment, which would require intensive and sustained treatment). Now, if Additivity were true, it would follow that, not only am I causally responsible for the individual deaths of each of those people whose lives I could have saved, but I am also causally responsible for the sum of those deaths. That sum includes the deaths of the "expensive" victims as well as the deaths of others who I couldn't have saved if I had saved any of the expensive victims.

In other words, on the basis of Additivity, we could truly say things like this: "One by one, picture everyone in the world whose lives you could have saved in some way or other (presumably, this is going to be many, many lives). Now picture all of them together. It turns out you are causally responsible for the collective outcome consisting in all of those people's deaths."<sup>19</sup>

Someone will surely ask: But, isn't it bad enough that we have to say that I caused all those deaths individually? In particular, isn't it bad enough that we have to say that I caused the death of a very "expensive" person, and also the deaths of all the other persons that I couldn't have saved if I had saved him? Why is it any worse to say that I caused them collectively?

In order to answer this question, let us look at an analogy. My hope is that this analogy will help bring out why we should regard the failure of the Additivity principle as a relief (or, at any rate, as some relief).<sup>20</sup> Failures of additivity are particularly evident when it comes to harms across different individuals. Pains don't seem to "add up" across people. For instance, causing 1 unit of pain to each of 10,000 people seems to be less bad than causing 10,000 units of pain to one single person. In other words, many small pains inflicted to different individuals don't seem quite as bad as a single excruciating pain inflicted to a single individual: pains are not "inter-personally additive."<sup>21</sup>

Perhaps similarly, then, causing the deaths of many starving children in Africa individually is not quite as bad, but is much less bad, than causing all of them collectively. When there is a failure of Additivity, the evils we cause individually don't add up to a single collective evil for which we are also causally responsible. The failure of Additivity is a relief, then, in that, although we are causally responsible for all those individual evils, we are not, in addition, causally responsible for the collective evil. In other words, although we are causally responsible for each of those deaths, we are not causally responsible, so to speak, for all that, or for that much, death.

## References

- Beebe, H. (2004) "Causing and Nothingness," in Collins, Hall, and Paul (eds.), Causation and Counterfactuals, Cambridge, Mass: MIT Press, 291-308.
- Hall, N. (2004) "Two Concepts of Causation," in Collins, Hall, and Paul (eds.), Causation And Counterfactuals, Cambridge, Mass: MIT Press, 225-276.
- Kim, J. (1993) "Causes and Counterfactuals," in Sosa and Tooley (eds.) Causation, Oxford University Press.
- Lewis, D. (1986) "Causation," reprinted with Postscripts in Philosophical Papers, vol. II, New York: Oxford U.P, 159-213.
- Lewis, D. (2004) "Causation as Influence," in Collins, Hall, and Paul (eds.), Causation and Counterfactuals, Cambridge, Mass: MIT Press, 75-106.
- McGrath, S. (2005) "Causation By Omission: A Dilemma," Philosophical Studies 123, 125-148.
- Parfit, D. (1978) "Innumerate Ethics," Philosophy and Public Affairs 7, 4, 285-301.
- Sartorio, C. (2004) "How To Be Responsible For Something Without Causing It," Philosophical Perspectives, Ethics, v.18, 1, 315-336.
- Sartorio, C. (2005) "A New Asymmetry Between Actions And Omissions," Noûs 39, 3, 460-482.
- Schaffer, J. (2005) "Contrastive Causation," Philosophical Review 114, 3.
- Singer, P. (1972) "Famine, Affluence, and Morality," Philosophy and Public Affairs 1, 3, 229-243.
- Taurek, J. (1977) "Should the Numbers Count?," Philosophy and Public Affairs 6, 4, 293-316.
- Thomson, J. (1977) Acts and Other Events, Ithaca: Cornell University Press.
- Thomson, J. (2003) "Causation: Omissions," Philosophy and Phenomenological Research 66, 1, 81-103.

Unger, P. (1996) Living High And Letting Die, New York: Oxford University Press.

---

\* Thanks to Juan Comesaña, Daniel Hausman, John Hawthorne, Russ Shafer-Landau, Alan Sidelle, Elliott Sober, and Dennis Stampe for helpful comments on earlier versions of this paper.

<sup>1</sup> Additivity presupposes that there are sums and that they enter in causal relations. Sums can be mereological sums, conjunctive facts, complex states of affairs, etc.

<sup>2</sup> This suggests a way of generalizing Additivity, which I have formulated in terms of an agent's acts, to all events. The generalized version would read: "If C1 causes E1 and C2 causes E2, then C1+C2 causes E1+E2." I have opted to focus on the version involving acts for reasons that will be apparent later (see section 4 below).

<sup>3</sup> A few people reject the possibility of causation by omission (see, e.g., Beebe (2004)). I will not attempt to defend the idea that omissions can be causes here. Instead, it will be an assumption of this paper.

<sup>4</sup> Naturally, If Additivity holds generally, then we should also expect it to hold for "mixed" cases involving both actions and omissions.

<sup>5</sup> Someone might argue that this depends on the contrast class with respect to which we are making the assertion: I didn't cause all the deaths rather than none of them (since I couldn't have saved all the soldiers), but I did cause all the deaths rather than only some (all but one) of them (given that I could have saved one soldier). However, there does seem to be a non-contrastive reading of "I caused all the deaths" (one according to which I caused all the deaths "simpliciter"), and (so far, at least) it is very plausible to think that the claim is false on this reading. (See also n.7 below.)

<sup>6</sup> A further reason to say that I caused each of the soldiers' deaths is this (thanks to Alan Sidelle for the suggestion). Imagine that I had saved the life of one soldier (say, Jim). Then it seems that

---

we would want to say, for each of the other soldiers (say, in particular, Bob), that I caused his death. But, if I cause Bob's death in the scenario where I save Jim, then, surely, I must also cause his death in the scenario where I save no one. (This argument assumes that we have a clearer intuition that I caused Bob's death in the case where I save someone else than in the case where I save no one. Assuming this is right, why is it? Maybe it's because, if I save someone else, the fact that I could have saved Bob becomes more salient; hence, it seems clearer that I was one of the causes of his death.)

<sup>7</sup> Both of these readings are non-contrastive. Someone might think of offering a contrastive explanation of the ambivalence we feel towards "I caused all the deaths." According to such an explanation, whereas I didn't cause all the deaths rather than none of them, I did cause all the deaths rather than all but one of them. But I think my explanation is better; in particular, it doesn't assume the contrastiveness of causation, which is highly controversial (for a defense of the idea that causation is contrastive, see Schaffer (2005)). (See also n.5 above.)

<sup>8</sup> What if I could have saved all the soldiers by packing my powerful shotgun the night before? If so I did cause the sum of the deaths (by failing to pack my shotgun). But I am not denying this. What I am denying is that I caused the sum in virtue of my compound omission consisting in my failing to act today in any of the ways that would have resulted in a soldier being saved. This is what rejecting Additivity amounts to.

<sup>9</sup> Assuming that an asymmetry of this type between actions and omissions exists, why does it exist? My guess is that it is probably a consequence of a more general asymmetry between causation by omission and causation by commission. In particular, I think it has to do with the fact that causation by omission is in general more sensitive to what happens in counterfactual scenarios than causation by commission is. (As an illustration of this phenomenon, compare: (a)

---

regardless of what would have happened if I had not kicked my victim, I caused his bruises given that I kicked him, and (b) if I couldn't have saved one of the soldiers by using my gun, then I didn't cause his death by failing to use it; for discussion of this asymmetry see Sartorio (2005).)

<sup>10</sup> Note that the claim is not that counterfactual dependence is necessary for causation. Almost everyone believes that this is false. See Lewis (1986).

<sup>11</sup> Other alleged counterexamples that have been offered in the literature are: “double prevention” cases (Hall (2004)), “accidental” dependencies (Sartorio (2005)), and “cheap” dependencies (such as failures to water plants that no one expected to be watered; see Thomson (2003) and McGrath (2005)).

<sup>12</sup> Here is David Lewis: “When asked to suppose counterfactually that C does not occur, we don't really look for the very closest possible world where C's conditions of occurrence are not quite satisfied. Rather, we imagine that C is completely and cleanly excised from history, leaving behind no fragment or approximation of itself” (Lewis (2004), p. 90). See also Lewis (1986), p. 211.

Could it be argued that, to the extent that, say, the assassin's shooting with another gun would have been a different event, the ensuing death would also have been a different event? There is no reason to think that this will be true in general: the death could have occurred in basically the same way even if the assassin had shot in a different way and with a different gun.

<sup>13</sup> Pace Thomson, who writes: “I am inclined to think, in fact, that if an event C causes an event E, then if we combine C with C' to form a compound event, then... the compound event also causes E... In particular, the event consisting in Sirhan's retracting the trigger the first one-hundredth of an inch and the rest of his shooting of [Robert] Kennedy caused Kennedy's death” (Thomson (1977), p.65). I agree that it is plausible to say that the sum of all the parts of Sirhan's

---

shooting of Kennedy is a cause of Kennedy's death (after all, they are all part of the shooting of Kennedy); however, it is much less plausible to say that the sum of Sirhan's shooting of Kennedy and the Pope's saying a prayer is a cause of Kennedy's death.

<sup>14</sup> For Lewis, the fact that removing the sum is removing the parts has the following advantage. In a case of symmetric overdetermination, say, one where two bullets simultaneously pierce a person's heart, the death doesn't counterfactually depend on either shooting but it does counterfactually depend on the sum of the two shootings. Hence, although we cannot say that the individual shootings caused the death, we can at least say that their sum did (Lewis (1986), p. 212). However, what Lewis failed to point out is that the sum of the two shootings and any other event in the world also comes out as a cause of the death. This seems wrong. Again, this is a reason to believe that, even if there is counterfactual dependence between the sum and the death, there isn't causation.

<sup>15</sup> I can still be blamed for saving one of the soldiers instead of another, say, if the one I saved is single and the other had a family. But I cannot be blamed for the fact that more than one death occurred.

<sup>16</sup> See, e.g., Beebee (2004).

<sup>17</sup> Notably, Singer (1972) and Unger (1996).

<sup>18</sup> Plausible as it might seem, I think that this explanation would require tinkering. In Sartorio (2004) I argue that the inference from the absence of causal responsibility to the absence of moral responsibility is not that simple, and I propose an alternative picture of the relation between causal and moral responsibility (one in which we can still draw lessons about moral responsibility from attributions of causal responsibility, but in a less straightforward way).

---

<sup>19</sup> Of course, it wouldn't follow from this that each of us is morally responsible for the collective outcome consisting in the sum of those people's deaths. However, again, it would be bad enough if it turned out that we are causally responsible for it. For, presumably, we feel that we are not morally responsible for it because we are not causally responsible for it. (Someone could try to offer this as an alternative explanation: we are not morally responsible for the sum of those deaths because we couldn't have prevented all of them at the same time. However, I don't think that this is really an alternative explanation. Why does the fact that we couldn't have prevented all the deaths explain the fact that we are not morally responsible for the sum? A natural answer is: because, given that we couldn't have prevented all the deaths, we didn't make the sum happen, and, if we didn't make it happen, we cannot be held responsible for it.)

<sup>20</sup> To be clear: I don't mean to suggest that the failure of Additivity is a relief with regard to the old problem of famine (the fact that we seem to be causes of the individual deaths of people in distant places). Rather, it's a relief in the sense that it helps avoid a new problem (the problem that would exist if we were causes of all those deaths taken collectively).

<sup>21</sup> See Taurek (1977) and Parfit (1978) for discussion. Perhaps pains are also not intra-personally additive either, when they occur at different times in a person's life. That is, several mild pains undergone at different times in a person's life might be less bad than a single excruciating pain at one single time.