

Causal Relevance, Permissible Omissions, and Famine Relief

Chad VANCE[†]

ABSTRACT

Failures are sometimes, but not always, causally relevant to events. For instance, the failure of the sprinkler was causally relevant to the house fire. However, the failure of the dam upstream to break (thus inundating the house with water) was not. Similarly, failures to prevent harms are sometimes, but not always, morally wrong. For instance, failing to save a nearby drowning child is morally wrong. Yet, you are also in some sense “allowing” someone on another continent to drown right now, and this seems permissible. Here, I argue that these two issues are connected. Roughly, I argue that it is *prima facie* morally wrong to fail to prevent a particular harm if and only if one’s omission is causally relevant to that harm’s occurrence. The result is that, contrary to what Peter Singer claims, failing to donate to famine relief is not morally equivalent to failing to rescue a drowning child in a shallow pond.

1. A dilemma for omissions

Intuitively, failures are sometimes causally relevant to events. Consider:

Faulty Sprinkler Some oily rags are piled near some exposed wire. The rags catch fire. A nearby sprinkler system, designed to activate under extreme heat, malfunctions and fails to put out the fire. The fire consumes the building.

However, sometimes failures “allow” events to occur in ways that do *not* seem to be causally relevant. Consider:

Unbroken Dam Some oily rags are piled near some exposed wire. The rags catch fire. Ten miles upstream there is a dam which, if it breaks, will release a flood of water that will inundate the building where the fire is occurring, putting the fire out. The dam fails to break, and the fire consumes the building.

While both the failure of the sprinkler system and the failure of the dam to break in some sense “allowed” the fire to occur, intuitively, we would only cite the failure in the first case (Faulty Sprinkler) as causally relevant to the event of the fire.

Thus, we have a causal dilemma: If failures are *never* causally relevant, then we obtain the counter-intuitive result that the failure of the sprinkler system *is not* causally relevant to the event of the fire. On the other hand, if failures are *always* causally relevant, then we obtain the counter-intuitive result that the failure of the dam to break *is* causally relevant to the event of the fire.¹

[†] Department of Philosophy, College of William & Mary; Email: cvance@wm.edu

¹ Sarah McGrath (2005, 144) states it this way: “DILEMMA: Either there is no causation by omission, or there is far more than common sense says there is. DILEMMA has it that either way, common sense is wrong.” The second horn of this dilemma is known as the problem of ‘Profligate Omissions’. The term seems to originate in Peter Menzies

But, there is an exactly parallel dilemma regarding the moral permissibility of omissions (i.e., failures to prevent harms). Consider: It is at least sometimes seriously morally wrong to allow a harm. For instance, failing to save a child who is drowning in a nearby shallow pond is seriously morally wrong. But, surely it is not *always* morally wrong to allow a harm. For instance, sadly, while you have been reading this, a child has probably drowned somewhere or other. In some sense you have “allowed” this to occur. After all, there is a metaphysically possible scenario where you prevent that child’s death. Yet, this failure on your part seems morally permissible.

Thus, we have a moral dilemma: If omissions are *never* morally wrong, then we obtain the unacceptable result that it is morally permissible to watch a child drown without intervening. On the other hand, if they are *always* wrong, then each of us is guilty of wrongly allowing the tens of thousands of deaths that occur each day due to causes which we could have easily prevented (e.g., drowning, auto accidents, starvation, and so on).² Ordinary intuition endorses the conclusion that, morally, failing to prevent a harm is *sometimes* wrong and *sometimes* permissible.

I believe it is no coincidence that these two dilemmas are parallel. Indeed, I will argue here that they are one and the same. For, I believe that *our moral principles have metaphysical underpinnings*; and, in this paper, I will argue that the moral wrongness of an omission is grounded in its causal relevance. Specifically, I will argue for the following claim:

The Wrongness of Omissions It is *prima facie* morally wrong for an agent (S) to fail to prevent a particular harm (H) iff S’s omission (O) is *causally relevant* to H.

My argument for this claim is quite simple. In §2, I will offer a metaphysical model, or analysis, of the causal relevance of failures with respect to events which combines two very popular theses: (a) a counterfactual account of causation, and (b) an account of modal proximity which privileges actual-world regularities. In §3, I examine this metaphysical model in the context of that variety of failures which is of particular interest to the ethicist—namely, failures to prevent harms (henceforth, ‘omissions’). I then address some potential objections in §§4-6 before applying my proposal to the issue of famine relief (in §§7-8). Roughly, I will argue that one’s failure to save a drowning child in a shallow pond is like the failure of the sprinkler in

(2004). See also Thomson (2003), McGrath (2005), and Bernstein (2013) for a discussion of this problem within the context of moral responsibility.

² Why is it true that I allow each and every one of the tens of thousands of deaths which I could have easily prevented? Consider this claim in the context of the problem of profligate omissions (see note 1, above). On a standard counterfactual analysis of causation, if it is true that, e.g., “If I had flown to Sydney on Monday at noon and pulled Suzie out of the water, then Suzie would not have died”, then I am a cause of Suzie’s death. Meanwhile, if it is true that, “If I had flown to New York on Monday at noon and pulled Timmy out of the water, then Timmy would not have died”, then I am *also* a cause of Timmy’s death.

Here, when I say that I have allowed both of these deaths, I mean only that I have failed to prevent that which I could have prevented (for, there is a possible world where I save Suzie, and a possible world where I save Timmy; though, in the actual world, I save neither of them). Of course, there is no possible world where I save *both* Timmy and Suzie. After all, I cannot simultaneously be in both Sydney and New York on Monday at noon. So, if it is *always* wrong to fail to prevent a harm, plausibly this wrongness is often overridden (for—following Ross, 1930—I take such wrongness to be of the *prima facie* variety, which can be overridden); namely, overridden by the fact that I could not possibly save everyone who needs saving due to my limited amount of time and resources, my inability to be in two places at once, and so on. For instance, the wrongness of failing to save Suzie (if such a failure were in fact wrong) would be overridden if I was, at the time, saving Timmy instead. Still, if omissions are always wrong, then I ought to at least have saved *one or the other*—yet, it seems permissible to save *neither* of them.

Faulty Sprinkler, while one's failure to donate to famine relief is like the failure of the dam to break in Unbroken Dam. Accordingly, by The Wrongness of Omissions, it will turn out that the former omission is *prima facie* morally wrong, while the latter omission is not.

2. Causal relevance and modal proximity

Consider the following Lewisian counterfactual analysis of causation (see, e.g., Lewis 1973):

The Causal Relevance of Failures The failure of some event (A) to occur is causally relevant to some particular event (E) if and only if:

- (i) $\neg A \rightarrow E$; i.e., the nearest possible worlds where A fails to occur are ones where E occurs, and
- (ii) $A \rightarrow \neg E$; i.e., the nearest possible worlds where A occurs are ones where E does not occur.

On this analysis, the failure of the sprinkler in Faulty Sprinkler is causally relevant to the fire (for, if the sprinkler fails, then the fire occurs; and if the sprinkler does not fail, then the fire does not occur). Unfortunately, this analysis entails that the failure of the dam to break in Unbroken Dam is *also* causally relevant to (or, as Lewis would have it, is *a cause of*)³ the event of the house fire. This seems mistaken.

Lewis handles this problem by appealing to context. On his view, we simply ignore the failure of the dam to break when determining causal relevance because, within that context, it is deemed irrelevant (though, strictly speaking, it was just as much a cause of the fire as the failure of the sprinkler). But, I wish to do better than this. I believe we can capture the intuition behind our contextual discrimination by adding a third condition to the two stated above—one which appeals to modal proximity, as follows:

- (iii) There are nearby possible worlds where A occurs.

Though the fire occurs if the dam fails to break, and *does not* occur if the dam *does* break, there are no nearby possible worlds where the dam breaks. It is metaphysically *possible* that the dam could have broken, to be sure—but the worlds in which this occurs are remote with respect to our world; for (assuming that the dam was not on the verge of collapse when the house fire occurred), much about our actual world would have to change in order for this to happen. Meanwhile, there *are* nearby possible worlds where the sprinkler goes off and douses the fire.⁴ In short, with the addition of criterion (iii), it turns out that the failure of the sprinkler *is* causally relevant to the occurrence of the house fire, while the failure of the dam to break *is not*. My proposed amendment, then, has some initial intuitive plausibility.

Before moving on, however, a metric for determining the 'nearness' of possible worlds is required. Now, how near a world is to the actual world is sometimes proposed to simply be a function of *spatiotemporal overlap*; i.e., the fewer states of affairs a world shares in common

³ Throughout this paper I will be concerned with *causal relevance*, rather than *causation*. Contrary to what Lewis says (2000, 195), plausibly, failures or absences cannot *cause* events (this is controversial). However, surely they may nevertheless be *causally relevant* to events (this is less controversial), and the latter is all that I require.

⁴ If there are any doubts about this claim, see note 14 below, where I attempt to assuage them.

with the actual world, the further that world is from actuality. However, in light of certain counterexamples, it is more commonly thought that spatiotemporal overlap only determines modal proximity *within that subset of worlds where actual-world regularities are held fixed*. To understand why, consider this case:⁵

Nixon and the Button Consider the following three possible worlds:

- w^* – In the actual world, Richard Nixon never pushes “the button” which would initiate a global nuclear war. Furthermore, no such war ever occurs.
- w_2 – In world w_2 , Nixon pushes the button. A global nuclear war occurs.
- w_3 – In world w_3 , Nixon also pushes the button. But, a global nuclear war never occurs because the button malfunctions.

If we were to ask, ‘What would have happened if Nixon had pushed the button?’, the intuitive answer is: ‘A nuclear war would have occurred.’ However, if counterfactuals are analyzed in terms of the nearness of possible worlds, and the nearness of a possible world is determined solely by its spatiotemporal overlap with the actual world, then the answer is rather: ‘The button would have malfunctioned, and a nuclear war would not have occurred.’ For, in terms of differences of states of affairs, the world where the button malfunctions (w_3) is *much* nearer to the actual world than the world where a nuclear war occurs (w_2). After all, in w_2 , our planet became an uninhabitable war-torn wasteland roughly fifty years ago.

In order to obtain the verdict that it is really w_2 (rather than w_3) that is nearer to the actual world, we must give priority to *regularities* or *norms* rather than overlap. In the actual world, buttons regularly *work* when pressed, rather than malfunction. In other words, the failure of a button—especially such an important one!—is *irregular* (i.e., *abnormal*). The world where the button malfunctions contains a departure from this norm. Only if regularities are given priority when ordering the closeness of possible worlds are we able to preserve the intuition that w_2 is nearer to the actual world than w_3 . So, I suggest that we do so.⁶

3. Strong and weak omissions

Let ‘omission’ refer to a failure of a moral agent to prevent a harm. Omissions, then, are a subset of failures. As such, my proposed account of the Causal Relevance of Failures from the previous section may be applied to omissions as follows:

The Causal Relevance of Omissions The failure of an agent (S) to perform some action (A) is causally relevant to a particular harm (H) if and only if:

- (i) $\neg A \rightarrow H$; i.e., the nearest possible worlds where S fails to perform A are ones where H occurs, and

⁵ This popular case is generally credited to Kit Fine (1975).

⁶ There is precedence for the use of the notion of ‘normal’ in one’s analysis of omissions. For instance, Sarah McGrath (2005) does so, where it is ‘normal’ for x to ϕ just in case there is some standard according to which x is *supposed* to ϕ (where ‘supposed to’ is understood in terms of moral, biological, physical, social, and even accidental regularities). Peter Menzies (2004) also makes use of the notion, claiming that causation is to be understood in terms of models with parameters that represent the ‘normal’ conditions underlying a given event. Note also that we should be aware of further complications with the metaphysics here, which I have greatly simplified. For more details about cases in the style of Nixon and the Button, see Lewis (1979), Schaffer (2004).

- (ii) $A \rightarrow \neg H$; i.e., the nearest possible worlds where S performs A are ones where H does not occur, and
- (iii) S performs A in nearby possible worlds (where ‘nearby worlds’ refers to that subset of the metaphysical possibilities which privileges actual-world regularities).

My proposal that the metaphysical underpins the moral is initially promising. For, the account above—in conjunction with The Wrongness of Omissions from §1—delivers verdicts that align with our ethical intuitions in a number of cases. Consider, for instance:

Abe’s Daughter Abe does not feed his daughter, and she dies. Bea, who lives on another continent and has never heard of Abe or his child, also does not feed her.

On my analysis, Abe’s failure is *causally relevant* to the child’s death, while Bea’s is not. For, (i) if Abe fails to feed his daughter, she dies, (ii) if he feeds her, she does not die, and—most importantly—(iii) there are nearby worlds where he *does* feed her. Claim (iii) is true since (presumably) the norms and regularities of Abe’s world did not include any physical or epistemic barriers which would have prevented him from doing so. Meanwhile, though Bea *also* could have prevented the death, all of those worlds where she does so are remote. The result is that Abe is guilty of what I shall call a *strong omission* (i.e., a causally relevant omission), while Bea is only guilty of a *weak omission* (i.e., a causally non-relevant omission). Metaphorically speaking, Abe is the faulty sprinkler, while Bea is merely the unbroken dam. On my account, this entails that Abe *wrongfully* fails to prevent his child’s death, while Bea permissibly fails to do so.⁷ This verdict seems correct. Thus, my proposal has some initial plausibility.

What is more, if my causal/modal analysis of failures continues to map onto what are intuitively the correct moral verdicts in a variety of cases, then I take this to be evidence in support of my central claim: The morality of omissions has metaphysical underpinnings. It should not surprise us if this were true. For, it is a common sentiment that *ought implies can*.⁸ But, ‘can’ is a modal term. If we say that S *ought* to have prevented harm H (by performing A), then it had better be the case that S *could* have performed A—and, in possible worlds speak, I take this to mean simply that some possible world where S performs A is *nearby*. But, modal proximity comes in *degrees*—as such, so does ‘can’. Intuitively, if there are *many* nearby worlds where S performs A, then S *can* perform A in a very strong sense; and if there are few such worlds—or worse, none at all, such that all worlds where S performs A are remote—then S “can” perform A in only a very weak sense; and the view I defend here is simply that the strength of one’s obligation, or ‘ought’, is dependent upon the strength of one’s ‘can’.

4. First problem: permissible strong omissions

⁷ Strictly speaking, my view is that *every* omission is prima facie morally wrong, but that this wrongness comes in degrees; namely, in accordance with its causal relevance, which (I believe) also comes in degrees. (Otherwise, my proposal is open to the charge of moral vagueness; e.g., where is the boundary between a morally permissible omission and an impermissible one?) The result is that we do, I think, have very weak moral duties to refrain from what I have called ‘weak’ omissions (e.g., Bea’s omission). However, such duties are so easily overridden, that they are treated in this paper as negligible (in much the same way as the gravitational influence that, say, the Andromeda Galaxy exerts upon the Earth is so vanishingly small as to be negligible).

⁸ See, e.g., Vranas (2007).

Now: *Does* my metaphysical proposal map cleanly onto our moral intuitions in a variety of cases? It may seem that there are a number of instances where it does not. Consider, for instance, the following potentially problematic case:

Fred Poolside Fred is listening to music on his headphones and engrossed in a book poolside, while a child drowns nearby. Fred fails to see or hear the child. The child dies.

We may initially be inclined to conclude that Fred has *strongly* allowed the child to drown. After all, Fred is *right there* as the child dies. It would have been quite easy for Fred to have looked up, become aware of the child's urgent situation, and pulled him out of the water. These factors support the conclusion that there are nearby worlds where Fred prevents the drowning (in which case, he is guilty of a strong omission, and has therefore acted wrongly, which seems mistaken). In short, my proposal appears to entail that Fred has acted wrongly, though intuitively he has not. (In the next section, I will examine another sort of difficulty: Cases where my thesis appears to count *seemingly impermissible omissions as weak omissions*.)

Now, we might simply sidestep the problem at hand by suggesting that causal relevance is only a *necessary* but not a *sufficient* condition for the prima facie moral wrongness of an omission. We could then accept that Fred's omission *is* a strong omission, causally relevant to the child's death, though causal relevance alone is not sufficient for moral wrongness.⁹ Alternatively, we might insist that he *has* acted wrongly, but is not blameworthy for doing so. This might explain the likely aftermath: Fred feeling the weight of immense guilt, while others will try to console him and tell him not to blame himself for what happened. But, I think we can handle Fred's case without weakening my proposal in either of these ways.

Now, while the obstacles between Fred and the prevention of the child's death are very weak in one sense—namely, overcoming them would not require much *physical* effort—there is another sense in which they are strong. For, while Fred's relationship to the situation is *physically proximate*, it is *epistemically remote*—and this is causally relevant. Consider: It is *normal* to respond to *reasons*, and *abnormal* to act without them. At the beginning of the scenario, Fred has no knowledge of the child's situation. As such, he has no reasons to attempt a rescue. Furthermore, regularity has established that such accidents are rare. As such, we do *not* typically have reasons to constantly scan every pool we visit, looking for drowning children, nor every restaurant we dine at, checking for choking patrons, and so on. So, Fred has no (or very little) reason to scan the pool for drowning children. Finally, it is *normal* for one to remain in a state of reading and music-listening until one has some reason to abandon them (it is regular, in fact, for *most* activities to confer some sort of *inertia* to those who engage in them).

As I have argued above, when considering which worlds are nearby, we must hold fixed the norms and regularities of the actual world. Once we do this, we may plausibly assert that those worlds where Fred rescues the child are much further from the actual world than those where he does not, and that—for this reason—Fred can rightly be said to have only *weakly* allowed the child to drown.¹⁰ In short, the sort of difficulty identified in this section can plausibly be dealt

⁹ For instance, perhaps *intending* harm (or at least foreseeing with a high degree of probability that what one is doing will result in harm) is also a necessary condition for (prima facie) wrongness—in which case Fred is not guilty of any wrongdoing, though his omission is a strong one.

¹⁰ At worst, even if there is *a* nearby world where Fred “just happens” to look up, sees the drowning child, and saves him, surely, since Fred has no reason to do so, there will not be very *many* such worlds. Furthermore, surely none of these worlds will be *nearest* to the actual world. In light of this, we might wish to alter condition (iii) of The

with once we understand that modal proximity is contextually restricted by the conservation of pre-existing, actual-world regularities which have been established by precedence, norms, initial conditions, inertia, expectations, reasons, and so on.

5. Second problem: impermissible weak omissions

5.1 Over-determination

Consider the following, a Frankfurt-style variant of James Rachels' (1975) Bathtub case:

Bathtub⁺ Jones sneaks into the bathroom, planning to drown his six-year-old cousin. The child slips and begins to drown on his own. Unbeknownst to Jones, Johnson is watching him, and will act such that, if Jones *tries* to save his drowning cousin, Johnson will prevent him from doing so. Johnson never intervenes, for Jones sits at the edge of the tub and watches his cousin drown without trying to save him.

On the present proposal, it seems that Jones has *not* strongly allowed his cousin to die, since there was a physical obstacle standing between himself and the prevention of the harm (albeit one that he was unaware of); namely, *Johnson*. Johnson's presence makes it such that any world where Jones saves his cousin is a remote one, so it turns out that—contrary to intuition—Jones's failure to prevent his cousin's death is morally permissible.

Here we have a case where my proposal appears to count a seemingly impermissible omission as a weak omission. Now, in response, one might simply point out that Bathtub⁺ is just a case of over-determination, which is problematic for *any* account of causation. As such, it need not concern us, since it is not a problem that is specific to the present proposal. Alternatively, we might revise the proposal so that causal relevance is only a *sufficient* but not a *necessary* condition for moral wrongness with respect to failing to prevent a harm. I wish to respond differently.

In Bathtub⁺, I accept that it is *true* that Jones cannot rightly be said to have strongly allowed the harm since the world where Jones saves his cousin *is* a remote one (due to Johnson's presence). Therefore, strictly speaking, Jones *does* permissibly allow his cousin's death. (After all, remember: *Ought* implies *can*.) But, this is not the end of the story. I have said that causal relevance is a necessary condition for moral wrongness *with respect to failing to prevent a harm*—but I have not said that it is a necessary condition for *moral wrongness simpliciter*. There are *other* things besides allowing harm that may be morally wrong (*doing* harm, for instance). Consider: For Jones, the *expected* result of trying to save his cousin was that he would have succeeded in saving the child (since he was completely unaware of Johnson's presence). Therefore, Jones's failure to respond to the situation of his drowning cousin reveals a bankruptcy of his moral character. As such, it would be just to hold him accountable for his inadequate response.

Bathtub⁺ seems to me to just be an instance of *moral luck* (albeit a new variety of luck, perhaps). For, if Johnson *had not* been present, then Jones' omission *would have* been causally relevant to his cousin's death. Jones is in some sense morally "lucky" that Johnson was present, since this made it such that Jones is not guilty of failing to prevent a harm which he could have

Causal Relevance of Omissions proposal to instead say 'S performs A in *many* nearby worlds' or 'in the *nearest* possible world'—and this would not affect my central thesis.

prevented—rather, he is only guilty of failing to *try* to prevent a harm which he (mistakenly) *thought* he could have prevented. In short, because there *was* (unbeknownst to him, and totally outside of his control) an action-preventer (namely, Johnson), Jones is “only” guilty of exposing his own moral bankruptcy.¹¹

5.2 *Misanthropes and psychopaths*

Yet, even without the over-determining presence of Johnson, the same worry may arise. Let us suppose that Jones *wants* his cousin to die—say, because he will then gain a large inheritance. If this desire for the child to die is a very strong one, then it acts as an epistemic obstacle and makes it quite unlikely that Jones would save his cousin. In short, if Jones wants his cousin to die (as it seems he does), then there is no nearby possible world where Jones saves him from drowning—and therefore, he is only guilty of a (morally permissible) weak omission, rather than a strong one. Yet we are confident that Jones’s failure in this case is morally reprehensible.¹² Worse still, this worry is generalizable, such that misanthropes never act wrongly when they allow harms, so long as they have a strong desire for the harm to occur. (After all, if ‘ought’ implies ‘can’, and a strong misanthropic desire renders it such that an agent “can” prevent a harm in only a very weak sense, then—on my view—it is also true that this agent also “ought” to prevent the harm in only a very weak sense.)

My reply is as follows: Strictly speaking, perhaps Jones has *not* strongly allowed the death of the child (due to his evil desires which served as a causal barrier). However, he *has* strongly allowed himself to cultivate a bankruptcy of moral character—and *this* omission has ultimately resulted in the death of a child. Therefore, the location of Jones’s morally impermissible omission is actually located not at the time of the child’s death, but somewhere earlier in his life, at a time when he failed to curb his own sinister dispositions.¹³

Psychopaths cannot be handled in this way, however. The true psychopath, as I understand him, is one who has not merely allowed himself to *cultivate* a poor moral character or an evil desire—but is rather one who simply never had the ability to be concerned with the welfare of others in the first place. If this is true, then it may be the case that, strictly speaking, psychopaths are not often causally (and therefore morally) responsible for failing to prevent harms. Admittedly, this sort of moral monster is a difficult one to handle, and my conclusion may be unpalatable to some. But, keep in mind that this verdict is merely the inevitable result for those who take seriously the intuition that ‘ought implies can’. Furthermore, the idea that one’s own dispositions can affect the causal relevance of one’s omission with respect to some event is really not so counter-intuitive. Consider, for instance, the following pair of cases:

The Bridge Keepers You cross two bridges:

¹¹ Other aspects of what Jones does might have been wrong as well. For instance, in the original story as James Rachels tells it, Jones *intends* the harm to his cousin as a *means* to an end (i.e., in order to gain an inheritance); and, plausibly, intending harm as a means to an end is morally wrong.

¹² See McGrath (2005) and Bernstein (2013) for further discussion of a similar worry.

¹³ Bernstein (2013) replies similarly to this sort of case, noting, “The idea is that certain dispositions are under the control of the agent. If an outcome is the result of a harmful disposition that she failed to correct, then she is responsible for the outcome in virtue of failing to correct the disposition that led to the outcome.” Similarly, as Vranas (2007) points out in the context of ‘ought implies can’, a student *ought* to turn in the paper that is due now, even if she has not written one (and therefore *cannot* turn in a paper). For, though she cannot write it *now*, she could have written it yesterday.

Bridge 1 (Larry): Larry, whom most people describe as “a pretty nice guy,” is standing on bridge 1, daydreaming as you walk by.

Bridge 2 (Troll): A troll lives under bridge 2, and has eaten every passerby who has ever attempted to cross it. As you stroll across the bridge above, the troll has an enormous urge to eat you, but he is turning over a new leaf. He musters all of his mental strength, and denies himself the tasty snack of your human flesh. You walk by innocently, not knowing that you have just narrowly avoided an untimely death.

It is not implausible, I think, to say that the troll has *strongly* allowed you to cross the bridge; that is, his failure to prevent you from crossing was very causally relevant to your success. After all, it was as if a great freight train was headed toward you and you didn’t even know it (i.e., the train of the troll’s hunger), but the troll intervened to hold it back. Meanwhile, Larry has only “allowed” you to pass in the very weak sense. Put another way, there seem to be many nearby worlds where the troll harms you, but none where Larry does.

Now, if internal dispositions can make one’s omission *more* causally relevant to a given event (as the troll’s dispositions seem to do with respect to your successful crossing of the bridge), then it stands to reason that they can also make one’s omission *less* causally relevant. So, if psychopaths are agents (if we can call them agents at all) for whom it is truly and uncontrollably “natural” to fail to prevent harms, then it may simply be the case that their omissions are *always* causally remote with respect to any given harm—no more causally salient than the omission of a rock or a tree or a breeze, for whom it is normal to ignore all nearby harms.¹⁴

6. Third problem: ethical relativism

I will now discuss a criticism more worrisome than the previous two. Consider this case:

Cannibalistic Society There is a society where cannibalism is commonplace, and widely accepted by all. Every week, someone in the society is randomly selected, sacrificed, and eaten. Dana grew up in this society. This week, she sits idly by as her kinspeople select their next victim and carry out their weekly ritual.

¹⁴ A worry may have occurred to the reader here: Perhaps broken sprinklers are like psychopaths, such that, if the sprinkler in Faulty Sprinkler is truly broken, then there are *not* any nearby worlds where the sprinkler prevents the fire—in which case its failure to go off is (contrary to what I have stated) *not* causally relevant to the house fire. I should like to say two things in response: (1) Perhaps modal proximity should prioritize what is normal for *types* rather than *tokens*. For, even if it is not normal for *this* particular (broken) sprinkler to function, it is normal for *sprinklers* to function. With this modification, there *would be* nearby worlds where the sprinkler prevents the fire. (Note that this response would render the psychopath’s failure to prevent a harm to be causally—and therefore morally—relevant to that harm as well, however.) (2) It is normal for sprinklers to be tested regularly, so there will inevitably be nearby possible worlds where the faulty sprinkler was discovered to be faulty, and then repaired. Alternatively, even if the particular sprinkler in question were *irreparably* broken, there will still be nearby possible worlds where *some* functioning sprinkler or other puts out the fire; e.g., because the original sprinkler was tested, discovered to be faulty, and promptly replaced. (Note that, unlike the previous response, this response would *not* extend to psychopaths. For, it is *not* a normal practice to test a human being’s capacity for empathy, and then promptly replace them with another human being, should that capacity be found to be irreparably broken.)

Let us assume that if Dana fails to intervene, then the death occurs; and if she intervenes, it does not. So far, so good: Her omission is causally relevant to the harm according to criteria (i) and (ii) of my proposal. However, her omission *fails* to meet criterion (iii), for it seems that there is *not* a nearby possible world where she *does* in fact intervene. Therefore, according to my proposal, Dana's omission is not causally relevant to the harm, and is therefore morally permissible. Applying this worry more broadly, it seems that my proposal reduces to a kind of ethical relativism: For, if agent A is deeply immersed in a culture which believes that allowing harm H is morally permissible, then plausibly all of the worlds where A prevents H will be epistemically (and therefore causally) remote. As such, on my proposal, it seems that A permissibly allows H (even if, in *our* culture, we generally believe such an omission to be morally repugnant). This is undeniably problematic.

I shall like to say three things in response—though I expect that some readers will not be satisfied by these. First, given our treatment of psychopaths in the previous section, our verdict for Dana is at least somewhat plausible. In most situations, there are no nearby worlds where psychopaths prevent harms. For them, preventing harms is only a remote possibility. Furthermore, since *ought* implies *can*, it follows that psychopaths have no moral obligation to prevent harms (since they “can” only in a very weak sense). But, something similar can be said of Dana. Having grown up in a society where cannibalism is an unquestioned norm, psychologically she has no motivating *reasons* to prevent the weekly sacrifices. Likely, it has simply never occurred to her to question the practice. Though she may not be a *true* psychopath (i.e., physically incapable of empathy), she has been sculpted by her society into something *like* one with respect to cannibalism. Brought up in a culture that is saturated with the acceptance and promotion of this practice, she is in some sense brainwashed in such a way that she “can” intervene to prevent the weekly sacrifice in only a very weak sense. In short, Dana's failure to prevent the harm is no more causally relevant to that harm than the failure of a rock or a tree to prevent it. In *Bridge Keepers*, Larry's failure to prevent you from crossing the bridge was not causally relevant to your crossing, since the possibility of his preventing you from doing so was so (psychologically) remote. Something similar should be said of Dana here, with respect to the sacrifice of the victim (the prevention of which was, for her, also psychologically remote).

Certainly, some will be sympathetic to this verdict. For instance, my assessment of Dana merely follows that of Susan Wolf's assessment of JoJo, the sadistic dictator who knows nothing other than the actions of his sadistic father, of whom she writes:

In light of JoJo's heritage and upbringing—both of which he was powerless to control—it is dubious at best that he should be regarded as responsible for what he does. It is unclear whether anyone with a childhood such as his could have developed into anything but the twisted and perverse sort of person that he has become. (1987, 54)

Wolf would say that Dana is not responsible for failing to prevent the sacrifice.¹⁵ I wish to go further, claiming that her failure is not even *causally relevant* to that event. In short, her failure is more like that of the dam to break than it is like that of the sprinkler to sprinkle.

Second, there are other ways besides allowing harm in which one can act wrongly; for instance, *doing* harm is *prima facie* morally wrong. My proposal here is limited only to the

¹⁵ See also Coates and Swenson 2013, who argue that moral responsibility comes in degrees, and is a function of one's ability to recognize and be receptive to moral reasons.

morality of *omissions*. Therefore, even if Dana's *failure to intervene* turns out to be permissible, this is consistent with the claim that those *doing* the sacrificing (and perhaps also the eating) are acting wrongly. So, there is no danger of my proposal collapsing into full-blown ethical relativism. On the contrary, on my proposal, there will still be many actions which are morally wrong, despite the fact that everyone in Dana's society believes them to be permissible (e.g., killing and eating people).

My third reply addresses those readers who might insist that Dana's omission is morally *wrong*, but admit that (due to the considerations raised here) perhaps she is not *blameworthy*. This reader may replace all instances in this paper of the word 'wrong' with the word 'blameworthy'. My thesis would be significantly weakened—as would my verdict for famine relief in §9—but still important. Namely, my thesis would instead be that *blameworthiness* (rather than wrongness) has metaphysical underpinnings; and, when applied to famine relief, my proposal would entail that, while failing to save a drowning child in a nearby pond *is* morally blameworthy, failing to donate to famine relief *is not*. Let us turn to that application now.

7. An intuitive case against the shallow pond analogy

Peter Singer has famously argued that we are morally obligated to donate to famine relief (1972). If we can prevent something very bad from happening without sacrificing something of comparable moral significance, he says, then we ought to do so. We *can* prevent something very bad (namely, death due to starvation) without sacrificing something comparable (namely, by donating money to famine relief). Therefore, we ought to do so. But, Singer then goes on to give his famous Shallow Pond case (1972, 231; see also 2009, 3), and this is often presented as the basis for a separate, second argument, distinct from the first.¹⁶ This argument by analogy is as follows:

1. Failing to save a child who is drowning in a nearby pond is seriously morally wrong.
2. But, failing to save a child who is drowning in a nearby pond is morally on a par with failing to donate money to famine relief (in order to save a child who is dying of starvation in a distant nation).
3. Therefore, it is seriously morally wrong to fail to donate money to famine relief.

However, we can easily undermine the analogy claim of the second premise with a simple pair of intuitive cases. Here is the first:¹⁷

¹⁶ Boonin and Oddie present Singer's argument in this way (2010, 10), for example. This interpretation of Singer is further supported by the fact that Singer spends a considerable amount of time arguing that there are no *morally relevant differences* between the two cases. For instance, Singer argues that the differences in physical proximity between the two cases, or the number of others who are able to help, or the degree to which you are certain that your efforts will do any good—these differences, he says, make “no *moral* difference”. (1972, 231, emphasis mine; see also Singer 1997). But, to say that there is no moral difference between two cases is just to say that they are morally on a par. In a similar effort, Unger (1996) goes to great lengths to argue that there are no morally relevant differences between his ‘Sedan’ and ‘Envelope’ cases (which roughly correspond to Singer's shallow pond and famine relief cases). For further examples of philosophers who present Singer and Unger as offering an argument by analogy, see Ashford (2011, 26), Cullity (2004, 10), Schmidtz (2000), and Woollard (2015, chs. 7-8). See also Timmerman (2015).

¹⁷ The following case is obviously inspired by Peter Singer's Shallow Pond case (1972, 2009), but also by his Bugatti case (1999; 2009, 13), which is itself adapted from Peter Unger (1996, 136).

Pond Bob is walking near a shallow pond where there are no other pedestrians nearby, and notices a drowning child. Bob is wearing an expensive watch—which, let us assume, is valued at \$1,000, would be ruined if it were to get wet, and would take too long to remove before jumping into the water were he to attempt to rescue the child in time to save his life. Planning instead to sell the watch and donate the proceeds to famine relief where it will save the lives of five children in a distant nation from starvation, Bob continues his leisurely stroll past the pond as the drowning child dies. Later that day, he sells the watch and donates the proceeds, as planned.

Intuitively, Bob acts wrongly in this case. While it would be *better* to save five children rather than one, nevertheless, ordinary intuition delivers an anti-utilitarian verdict here: Bob *ought* to save the one rather than the five.¹⁸ Now, consider the second case:

Charity Bob plans to sell his watch, which is valued at \$1000, and donate the proceeds to one of two causes: He can either donate the money to a charity that will save one child in a distant nation from starvation, or he can donate it to a charity that will save five children in a distant nation from starvation (who, let us assume, due to local economic circumstances, can be saved for \$200 each, rather than \$1000). Bob ultimately decides on the latter charity, and saves the lives of five children.

It seems that what Bob does in this case is morally permissible (and, plausibly, if these are his *only* two choices, even obligatory).¹⁹ But, then, we have a discrepancy: Whereas it is morally wrong for Bob to save the five children instead of the one in Pond, it is morally permissible for him to do so in Charity. Thus, intuition indicates that there must be some morally relevant difference between the two cases.

Once again: If Bob saves five starving children rather than one drowning child (in Pond), he does something morally *wrong*. On the other hand, if Bob saves five starving children rather than one starving child (in Charity), what he does is *not* morally wrong. But, there is only one difference between these two cases: Namely, in the first case the one child (whom he fails to save) can be saved only by being pulled out of a nearby pond, while in the second case the one child (whom he fails to save) can be saved only by donating money to famine relief. Thus, contrary to what Singer claims, these two failures are not morally on a par. My argument has been as follows:

¹⁸ Some readers will disagree with this verdict, stating that Bob does *not* act wrongly in Pond, and we really *should* refrain from saving the life of someone nearby in urgent need whenever doing so would impose a cost on us that could be put to better use elsewhere. However, I suspect that those who are inclined to reject utilitarianism will want to resist such a claim. Singer himself even endorses the anti-utilitarian verdict in this type of case, at least publicly. For instance, in his Bugatti example (1999; 2009), he claims that Bob ought to sacrifice his Bugatti (a car valued at \$2 million) in order to save the life of *one* nearby child in urgent need. Of course, it is possible that Singer privately believes that what Bob *really* ought to do is save his Bugatti, allow the one child to die, and then sell the car in order to donate some or all of the proceeds in order to save hundreds or thousands of lives—in which case, we should expect him to privately believe something similar about Bob in the case I have just described. No matter, since presumably most readers will share the intuition that Bob really *should* save the one child (rather than the five) in Pond—and my thesis here is directed at these readers.

¹⁹ For a defense of the latter claim, see: Pummer 2016.

1. It is much worse, morally, for Bob to fail to save the one child in Pond than it is for him to fail to save the one child in Charity.
2. The only candidate which might serve as a morally relevant difference between Pond and Charity is this: In Pond, Bob fails to rescue the child from drowning in a nearby pond, while, in Charity, Bob fails to donate to famine relief in order to rescue the child from dying of starvation in a distant nation.
3. Therefore, failing to save a child from drowning in a nearby pond is not morally on a par with failing to donate money to famine relief in order to save a child in a distant nation from starvation.

If our intuitions in Pond and Charity are correct, then we are committed to the conclusion that there is a morally relevant difference between failing to save a drowning child and failing to donate to famine relief—in which case, we ought to reject Singer’s argument by analogy. What is more, the metaphysical account of omissions that I have defended in this paper delivers this same verdict while also identifying what the morally relevant difference in question *is*. Let us turn to that topic now.

8. *Causal proximity and famine relief*

I have just argued that there is a morally relevant difference between failing to save a drowning child and failing to donate to famine relief. In the preceding sections, I supplied an account which identifies what that morally relevant difference *is*. It is this: Bob’s failure to save the one drowning child in the original Pond case is *causally relevant* to that child’s death. Meanwhile, his failure to save the five starving children in that scenario, or any of the six starving children in Charity, is *not* causally relevant to any of those deaths.

To understand why, it will be helpful to first point out that, on my analysis, causal relevance is strongly correlated with—but not identical to—physical proximity. A common response upon first encountering Singer’s Shallow Pond case is to deny that it is analogous to famine relief because the drowning child is nearby, while those who are dying of starvation are very far away. Because of this difference between the two cases (so the objection goes), it turns out that failing to save a drowning child is wrong, while failing to save a starving child is not. In short, physical proximity is morally relevant. Singer rejects this criticism—and, I think, rightly so.

A common way of testing for the moral relevance of proximity is to offer a case much like Singer’s Shallow Pond, but where the child in need is no longer nearby. For example, consider the following excellent case from David Boonin (2010):

E-velope You are checking out various vacation-planning web sites when a computer error mistakenly connects you to a web site sponsored by UNICEF. While you are reading the organization’s materials, a new bulletin is posted to the site. After reading the bulletin, you correctly believe that there is a particular child who has just contracted a fatal illness and who will die soon unless someone who is currently logged onto the site sends in \$100 to pay for the medicine she needs by clicking on an icon at the bottom of the screen, which will authorize a \$100 withdrawal from the user’s bank account. As you are considering whether or not to send in the needed money, a second bulletin appears. After reading this second bulletin, you correctly believe that the server UNICEF is using

is experiencing problems and that no one else will be able to log on to the site until it is too late to save this particular child. Finally, you look at a window at the bottom of the screen, after which you correctly believe that you are the only person who is currently logged onto the site. You correctly believe, in short, that this particular child will die if you do not click on the ‘send \$100’ icon and that this child will live if you do click on the ‘send \$100’ icon. You decide not to click the icon and turn off your computer, and, instead of living for many more years, this particular child dies.

We correctly judge that it would be morally wrong to refrain from saving the child’s life in this case, even though the child in need is very far away. Thus, physical proximity is not morally relevant. As Peter Unger says, “unlike many physical forces, the strength of a moral force doesn’t diminish with distance.” (1996, 33)

I will not take issue with this claim. I agree that it is morally wrong to allow the child to die in the E-velope case, and I agree more generally with the conclusion that physical proximity is not morally relevant. The mistake which drives the intuition behind the proximity objection is in thinking that it is *physical proximity* that is morally relevant, when in fact it is, rather, *causal proximity* that is morally relevant. While it may be true that the strength of our obligation does not diminish with *physical* distance, it *does* diminish with *causal* “distance” (i.e., *causal relevance*, as I have argued). What is more, the two *are* strongly correlated. That is, the further away some harm is in space (or time), the less causally relevant one’s failure to prevent that harm typically is. (Consider, for instance, Abe and Bea in Abe’s Daughter from §3.)

E-Velope is unusual in this sense. There, the omission is *causally proximate* (i.e., *causally relevant*) to the child’s death, despite being *physically remote*. (Note that physical and causal proximity can come apart in the other direction as well. For instance, recall that poolside Fred was *physically proximate* but *causally remote* with respect to the child’s drowning.) Boonin carefully describes the situation so that the agent is both vividly aware of the child’s urgent need and uniquely capable of helping. These factors “bridge the gap” of the physical distance, so to speak; and it is this bridge (or *tether*, perhaps) which pulls the agent into the sphere of causal influence surrounding the child’s death.²⁰

But, then, E-Velope is importantly different from our actual situation with respect to famine. For, rarely is there ever such a “bridge” or “tether” which draws us into the causal sphere of influence surrounding any particular famine-related death in a distant nation. (See the Appendix below for a lengthy defense of this claim.) Admittedly, the causal gap between affluent individuals and poverty-related deaths *would* be diminished if famine relief organizations were structured in the following way:

Extreme Solution There is a charitable organization called ‘It’s All Your Fault’, which assigns one starving child to each affluent person in the world. Charity workers then contact each of these affluent persons to let them know that they have been selected to sponsor a starving child, and that this child will receive life-saving aid *only* if their sponsor donates the necessary funds. Otherwise, the child will die.

²⁰ Frances Kamm is, I think, making a similar claim when she states (2007, 12.II) that our understanding of ‘nearness’ is confused. It should be understood, she says, not merely in terms of the physical proximity of the agent to the person in need—but rather as a function of several different varieties of proximity (e.g., the agent to the person in need, or to the threat, or to the *means* of preventing the threat, etc.). On my view, such factors will all act as tethers between the agent and the harm, each of them (typically) contributing to *causal* proximity.

Here, potential sponsors have been purposefully *inserted* into the sphere of influence surrounding the various harms (albeit, without their consent). In this case, failure to donate is—according to my proposal—causally relevant to a particular child’s death, and therefore morally wrong, in accordance with our intuition. Setting up a charity in such a way would of course be morally objectionable, even if it would elicit a greater number of donations and save more lives. Nevertheless, my point here is that Extreme Solution (like E-Velope) does not accurately represent the situation that we find ourselves in, causally, with respect to famine. For, if you are living in an affluent nation, then, due to a combination of factors including their distance, their number, and the way that famine relief organizations are presently structured, each individual harm or death is (typically) not only *physically*, but also *causally* remote. As Peter Unger puts it (1996, 48), our relationship with respect to harms in cases like Extreme Solution, E-Velope, and Singer’s Shallow Pond is “causally focused”, whereas, in the typical famine relief case it is “causally amorphous.”²¹

To sum up, the implications of my proposal for our test cases are these: Bob’s failure to save the drowning child in Pond constitutes what I have called a (prima facie morally wrong) *strong* omission, while his failure to save any of the other (starving) children in Pond or Charity does not. The result is that it is much worse, morally, for Bob to fail to save the one drowning child than it is for him to fail to donate money to save any of the other children. Applying this verdict to Singer’s argument by analogy, we see that it is much worse, morally, to fail to save a child drowning in a shallow pond than it is to fail to donate to famine relief. Thus, contrary to what Singer claims, the two cases are not morally analogous. More generally, the failure of an affluent individual to donate to famine relief is not a strong omission, and is therefore not (prima facie) morally wrong—at least, not on the grounds that one allows a particular child to die. (It might still be wrong for some other reason.)

Before concluding, however, I should note that I have *not* claimed the following: First, I have *not* said that donating to famine relief would be a *bad* thing. Clearly, it would be a very *good* thing. Second, I have not argued that strongly allowing a harm is always *all-things-considered* morally wrong. For there are obvious counter-examples to that claim:

Crowded Cliff You are standing at the edge of a cliff with a crowd of people. The crowd surges, and the stranger next to you begins to slip over the edge. There is no longer any room on the cliff for this person to stand. Just before they plummet to their death, however, you shove another person off of the edge of the cliff in order to make room for the slipping stranger, hoisting them back onto the ledge and saving their life.²²

If you had allowed the first stranger to fall you would have *strongly* allowed their death, to be sure. Yet, it seems that, morally, you ought to have let the stranger fall. For, I have only argued that strong omissions are *prima facie* morally wrong—and such wrongness can be overridden. In Crowded Cliff, the wrongness is overridden because the only way to prevent the death of the one

²¹ Of course, Unger mentions this distinction only in order to deny its moral relevance. To this end, he offers a case that is similar to Extreme Solution and claims that the ordinary intuition is that one does *not* act wrongly by failing to donate in such a case. I disagree. I suspect that most will have the intuition that one *does* act wrongly if they fail to donate in that case. Though admittedly, if I am wrong about this, then this would significantly undermine my central thesis.

²² This case is loosely inspired by Ross’s discussion of the duties of beneficence vs. non-maleficence (1930).

stranger is by causing the death of another. (For this reason, I have largely restricted my discussion to life-or-death situations—i.e., ones where someone either lives or someone dies—rather than what might be called ‘death-or-death’ situations.)

Finally, I am *not* claiming that a moral obligation to donate to famine relief could not be argued for in some other way. My claim has merely been that failure to donate is not morally wrong *on the grounds that it fails to prevent a particular harm*. There might be other reasons for why such inaction is morally wrong, however. For instance, a philosopher with utilitarian inclinations will almost certainly conclude that failure to donate to famine relief is morally wrong—not because such a failure allows some *particular* harm to occur, but rather because failure to donate fails to maximize the total utility. Alternatively, if it turns out that we are responsible for contributing to certain social institutions that are (in part) *causing* the harm to people in impoverished nations (see, e.g., Pogge 2002), then it is probably the case that we ought to donate to famine relief, since it is plausible that we are obligated to pay compensation for harms that we have caused.

In conclusion, my proposal entails that the failure of a moral agent to save a drowning child is, causally, like the failure of the sprinkler to put out the fire in Faulty Sprinkler. Meanwhile, the failure to donate to famine relief is like the failure of the dam to break in Unbroken Dam. In both pairs of cases there is a *causal* difference. But, when the failure in question is that of a moral agent to prevent a harm, this causal difference also generates a *moral* difference (or so I have argued). There are, I have noted, ways of “bridging the causal gap” between affluent individuals and poverty-related deaths (as in Extreme Solution); and, the closer we get to that ideal, the stronger our prima facie duty to prevent such deaths will be. But, so long as this morally relevant difference remains, Singer’s claim that the two cases are morally analogous remains mistaken, and a successful argument for the conclusion that we are morally obligated to donate to famine relief must be achieved in some other way.

Appendix: famine relief and causal diffusion

Here, I offer a more careful defense of the claim that our actual situation with respect to famine relief organizations suffers from what I will call ‘causal diffusion’. To illustrate what is meant by this, consider the following progression:²³

Life Preserver 1 (LP1) You stand alone at the shallow pond, where one child drowns, and may insert \$200 into a machine that will drop one life preserver into the pond near the child.

LP2 You stand alone at the shallow pond, where ten children drown, and may insert \$200 into a machine that will drop one life preserver into the pond near the children. However, there is only one life preserver left in the machine.

LP3 You and four other strangers stand at the shallow pond, where ten children drown. You each have \$200 in your pockets, and may insert that money into a machine

²³ These cases are loosely inspired by a progression of cases found in Kamm (2007, 348), as well as Unger’s distinction between what he calls ‘causally focused aid’ and ‘causally amorphous aid’ (1996, 48).

containing five life preservers, which will drop one life preserver into the pond near the children every time it receives \$200.

LP4 You and four other strangers stand at the shallow pond, where ten children drown. There is a machine containing five life preservers, which will only make one drop of some number of life preservers into the pond before malfunctioning. You and the other four strangers (who each have \$200) must pool your money, making one large deposit collectively, and then the machine will deposit up to five life preservers into the pond (one for every \$200 deposited).

LP5 You and three billion other affluent strangers stand at the shallow pond, where one billion children drown. There is a machine containing an unlimited number of life preservers, which will only make one drop of life preservers into the pond before malfunctioning. You and the other strangers must pool your money, making one large deposit collectively, and then the machine will deposit some number of life preservers into the pond (one for every \$200 deposited).

LP6 You and three billion other affluent strangers live in a world where there is a shallow pond on another continent, where one billion children drown. You and the other strangers may send money whenever you'd like to an organization overseas. The organization collects the money, and then distributes various types of aid however and whenever they see fit. Approximately 10% of the donations are used for expenses such as employee's salaries and advertising. While some of their projects are more/less cost-effective than others, on average the cost of saving a life is roughly \$200.

Now, to reiterate, my central proposal has been as follows:

The Causal Relevance of Omissions The failure of an agent (S) to perform some action (A) is causally relevant to a particular harm (H) if and only if:

- (i) $\neg A \rightarrow H$; i.e., the nearest possible worlds where S fails to perform A are ones where H occurs, and
- (ii) $A \rightarrow \neg H$; i.e., the nearest possible worlds where S performs A are ones where H does not occur, and
- (iii) S performs A in nearby possible worlds (where 'nearby worlds' refers to that subset of the metaphysical possibilities which privilege actual world norms and regularities).

With this account in mind, let us consider each of the six Life Preserver cases in turn: In Life Preserver 1, there is a one-to-one correspondence between the agent and the harm. In short, it is true that (i) the nearest possible worlds where you *do not* perform A (i.e., the act of spending the \$200) are ones where the particular harm (H) occurs, and (ii) the nearest possible worlds where you *do* perform A are ones where H is prevented. In short, both conditions (i) and (ii) of The Causal Relevance of Omissions are met.

In LP2, there is no longer a one-to-one correspondence, however. Now, condition (i) *is* met. For, the nearest possible worlds where you fail to perform A are ones where each of the ten particular harms (H_1, \dots, H_{10}) *does* occur. However, for any particular harm (say H_1), only 10% of the nearest possible worlds where you perform A are ones where H *does not* occur. This is

because the worlds where you perform A and H_1 is prevented are just as near to the actual world as the worlds where you perform A and, say, H_6 is prevented (or H_7 , or H_8 , etc.). Thus, if you fail to put money into the machine in case 2, we may correctly conclude that you have strongly allowed one of the ten children to die—though there is no fact of the matter about *which* child you have allowed to die, such that condition (ii) is only weakly met with respect to any particular harm.²⁴

In LP3, *both* conditions (i) and (ii) of my proposal are only weakly met. Regarding (i): In the worlds where you fail to perform A, all that we can say is that six children died rather than five (since, assuming the other four individuals pay up in all nearby worlds, each of the worlds where some set of six children die are equally nearby, such that there is no fact of the matter about which particular harm you have allowed). Regarding (ii): This criterion is only weakly met for similar reasons—since, e.g., the world where you and the others donate and $H_1 - H_5$ are prevented are just as near to the actual world as those where you and the others donate and, say, $H_2 - H_6$ (or $H_3 - H_7$, etc.) are prevented.²⁵

It should be clear that neither criteria (i) or (ii) are strongly met in any of the remaining cases. What is worth noting in LP4 is that it is the first case where there would not even be a fact of the matter about which life you have saved if you were to *actually* donate. (Since the money is pooled into a collective sum before being inserted into the machine, there is no fact of the matter about which particular life preserver is dropped into the pond as a result of *your* payment.) In that case, if you donate with the others, all that you may conclude is that there are five life preservers being dropped into the pond instead of four—end of story.

I have included LP5 for the following reason: While in LP4, you may rightly claim that you have decreased each child's odds of drowning from 60% to 50% (assuming that all five of you pay the machine), in LP5 a donation to the machine decreases the odds that any particular child will die by only a negligible amount.

Finally, in LP6, criteria (i) and (ii) are no longer even *weakly* met. Rather, they are not met at all. For, in LP6, some of the nearby possible worlds where you *fail* to donate will be ones where *more* lives are saved, compared to some of the nearby worlds where you *do* donate, in which (sometimes) *fewer* lives are saved (for instance, due to the various possible types of aid that the organization might choose to distribute, which vary in their cost-effectiveness).

It should be clear that, in Life Preserver 1, by failing to pay \$200, one is uniquely causally responsible for strongly allowing a particular harm. By LP6, we can no longer make sense of this claim, however. This is because, between each stage and its successor, there is an increasing amount of *causal diffusion*. It should also now be apparent that Singer's original Shallow Pond case (as well as E-Velope and Extreme Solution) are most similar to Life Preserver 1 in terms of their causal relevance. Our situation with respect to charitable famine relief organizations, however, is most similar, causally, to Life Preserver 6. In short, given the way that charities are presently structured, our present situation with respect to famine relief is such that there is no

²⁴ Of course, if you *do* pay the machine, then one of the ten possibilities will be actualized and it will become obvious that one particular child's life has been saved as a result of your action, A.

²⁵ Though, again, if you *do* pay the machine, then one of the various possibilities will be actualized and one particular child's life will be saved by your particular action—namely, whichever one grabs *your* life preserver—and you may correctly conclude that your action has saved that particular child's life.

particular harm (H₁, H₂, H₃, etc.) for which (i) H will occur if we *do not* donate, and (ii) H will not occur if we *do* donate.*

REFERENCES

- ASHFORD, E. 2011, "Obligations of Justice and Beneficence to Aid the Severely Poor", in: P. Illingworth, T. Pogge, and L. Wenar, eds., *Giving Well: The Ethics of Philanthropy*, Oxford: Oxford University Press, pp. 26-45.
- BERNSTEIN, S. 2013, 'Omissions as Possibilities', *Philosophical Studies* **167**, pp. 1-23.
- BOONIN, D. 2010, "Famine, Affluence, and Mortality", *A Priori: The Erskine Lectures in Philosophy, University of Canterbury*, **4**.
- BOONIN, D. and ODDIE, G. 2010, "Introduction: What's Wrong With Arguing?", in D. Boonin and G. Oddie, eds., *What's Wrong? Applied Ethicists and Their Critics*, Oxford: Oxford University Press, pp. 1-30.
- COATES, J. D. and SWENSON, P. 2013, "Reasons-Responsiveness and Degrees of Responsibility", *Philosophical Studies* **165**, pp. 629-645.
- CULLITY, G. 2004, *The Moral Demands of Affluence*, Oxford: Oxford University Press.
- FINE, K. 1975, "Critical Notice of David Lewis's Counterfactuals", *Mind* **84**, pp. 451-458.
- KAMM, F. M. 2007, *Intricate Ethics: Rights, Responsibilities, and Permissible Harm*, Oxford: Oxford University Press.
- LEWIS, D. 1973, "Causation", *The Journal of Philosophy* **70**, pp. 556-567.
- LEWIS, D. 1979, "Counterfactual Dependence and Time's Arrow", *Noûs* **13**, pp. 455-476.
- LEWIS, D. 2000, "Causation as Influence", *The Journal of Philosophy* **97**, pp. 182-197.
- MCGRATH, S. 2005, "Causation by Omission: A Dilemma", *Philosophical Studies* **123**, pp. 125-148.
- MENZIES, P. 2004, "Difference-Making in Context", in: L. A. Paul, E. J. Hall, and J. Collins, eds., *Causation and Counterfactuals*, Cambridge, MA: MIT Press, pp. 139-180.
- POGGE, T. 2002, "Responsibilities for Poverty-Related Ill Health", *Ethics & International Affairs* **16**, pp. 71-79.

* Many thanks to Chris Heathwood, who sparked my interest in this topic all those years ago, at the beginning of my philosophical career; to Duncan Purves for helpful comments on an early version of this paper; and to the editor and anonymous referees at *dialectica*.

- PUMMER, T. 2016, "Whether and Where to Give", *Philosophy & Public Affairs* **44**, pp. 77-95.
- RACHELS, J. 1975, "Active and Passive Euthanasia", *The New England Journal of Medicine* **292**, pp. 78-80.
- ROSS, W. D. 1930, *The Right and the Good*, Oxford: Clarendon Press.
- SCHAFFER, J. 2004, "Counterfactuals, Causal Independence, and Conceptual Circularity", *Analysis* **64**, pp. 299-309.
- SCHMIDTZ, D. 2000, "Islands in a Sea of Obligation: Limits to the Duty of Rescue", *Law and Philosophy* **19**, pp. 683-705.
- SINGER, P. 1972, "Famine, Affluence, and Morality", *Philosophy & Public Affairs* **1**, pp. 229-243.
- SINGER, P. 1997, "The Drowning Child and the Expanding Circle", *New Internationalist Magazine*, April 5.
- SINGER, P. 1999, "The Singer Solution to World Poverty", *The New York Times Magazine*, September 5, pp. 60-63.
- SINGER, P. 2009, *The Life You Can Save*, New York: Random House.
- THOMSON, J. J. 2003, "Causation: Omissions", *Philosophy and Phenomenological Research* **66**, pp. 81-103.
- TIMMERMAN, T. 2015, "There is Nothing Wrong With Letting a Child Drown", *Analysis* **75**, pp. 204-212.
- UNGER, P. 1996, *Living High and Letting Die: Our Illusion of Innocence*, Oxford: Oxford University Press.
- VRANAS, P. 2007, "I Ought, Therefore I Can", *Philosophical Studies* **136**, pp. 167- 216.
- WOLF, S. 1987, "Sanity and the Metaphysics of Responsibility", in F. Schoeman, ed., *Responsibility, Character, and Emotions: New Essays in Moral Psychology*, Cambridge: Cambridge University Press, pp. 46-62.
- WOOLLARD, F. 2015, *Doing and Allowing Harm*, Oxford: Oxford University Press.