

11 Causing and Nothingness

Helen Beebee

1 Introduction

According to the way we normally speak about the world, absences can be causes. We might say, for instance, that Jones's failure to close the fire doors was a cause of the raging fire that destroyed the building, or that Smith's failure to water her office plants was a cause of their death, or that lack of rain was a cause of the bush fire. So it seems that an adequate theory of causation is going to have to rule that there is such a thing as causation by absence.

I don't believe there is any such thing as causation by absence. In this essay I'm going to try to defend that claim against the objection just raised, partly by arguing that there are features of commonsense assertions and denials of causation by absence to which *no* theory of the metaphysics of causation ought to be doing justice, and partly by making a positive claim about the role absences play in our explanatory practices, a claim that, I think, allows me to rebut the objection that the repudiation of causation by absence flies in the face of too much of our commonsense understanding of causation to be taken seriously.

The reason I deny that there's any such thing as causation by absence is that I want to uphold the view that causation is a relation between events. To be rather more picturesque, I subscribe to what Helen Steward has recently dubbed the "network model of causation": The complete causal history of the universe can be represented by a sort of vast and mind-bogglingly complex "neuron diagram" of the kind commonly found in discussions of David Lewis, where the nodes represent events and the arrows between them represent causal relations.¹ Or, to put it another way, the causal history of any event is, as Lewis puts it, "a relational structure" (1986c, p. 216).

I think it's fair to say that the network model is the dominant model of causation in contemporary metaphysics—largely owing to the influence of Lewis and Davidson.² The network model lies behind most of the theories of causation currently on offer; it also lies behind a vast amount of the literature on the philosophy of mind.

If there *is* causation by absence, then the network model can't be right—or at least, it can't be right assuming there are no such things as negative events. And I assume in what follows that there are no negative events—which is to say, more or less, there are no events whose essence is the absence of a property or particular. If Jones's failure to close the fire door is not an event, and if this failure was a cause of the fire, then the full causal history of the fire is not exhausted by the network of

From J. Collins, N. Hall, and L. A. Paul, eds. *Causation and Counterfactuals* (Cambridge, MA: MIT Press, 2004): 291–308.

events and causal relations between them, for there will be no event of Jones's failure, and hence no causal relation between his failure and the fire. The network model cannot accommodate the fact—if it is a fact—that Jones's failure caused the fire, and hence cannot be the whole causal truth about reality.

How should we solve the problem? One solution—the one I favor—is to hang onto the network model of causation and deny that there is any causation by absence. On this view, Smith's failure to water her office plants was not a cause of their death. This seems to be a deeply unpopular solution—and surprisingly so, given the prevalence of the network model in contemporary metaphysics. Here's what Lewis has to say about it in "Void and Object" (chapter 10, this volume): "We could deny, in the face of compelling evidence to the contrary, that absences ever cause anything. We could deny, for instance, that the void is deadly. . . . Simply to state this response is to complete the *reductio* against it."³

The other obvious solution is to hang onto causation by absence and scrap the network model. This is the solution for which Hugh Mellor has long been arguing,⁴ and has been suggested more recently by Lewis in "Void and Object." According to Mellor's view, causation is not a relation—or at least, the most basic kind of causation is not relational. Rather, causal facts have the form "*E* because *C*," where *C* and *E* are facts and "because" is a sentential connective. Since facts are not particulars, on Mellor's view causation is not a relation. And since facts can be facts about absences, causation by absence causes no special problems on his view. In much the same vein, Lewis proposes that causation be analyzed in terms of counterfactual dependence not between events, but between propositions—propositions that need not assert that some event or other occurs.

A third solution—one that is suggested but not condoned by Lewis in "Void and Object"—is to take the relational kind of causation as the most basic kind of causation, and define causation by absence in terms of relational causation.⁵ According to this kind of solution, the network model captures all the *basic* causal facts, but there are other causal facts whose obtaining supervenes on the actual pattern of causal relations together with some extra counterfactuals.

I shall call the first view the *relationist* view, since it upholds the thesis that causation is a relation between events. I call the other two views *nonrelationist* views, since they both uphold the thesis that causation is not always a relation. The central claim of this essay is that, if our aim is to do as much justice as possible to commonsense intuitions about causation by absence, then it doesn't much matter whether we uphold relationism or nonrelationism.

This might seem like a surprising claim, since relationism, unlike nonrelationism, is forced to deny a whole host of causal claims that common sense holds true. As I

said at the beginning, undoubtedly many of us think that Jones's failure to close the fire doors was a cause of the fire, or that the lack of rain was a cause of the bush fire, and so on. Fair point: The relationist view really does deny all of that. But this is to tell only half the story. For common sense also makes a lot of *negative* causal claims about absences. Take Brown, who lives on the other side of the city and has no connection whatever with either Smith or Jones. Commonsense intuition has it, I think, that neither Brown's failure to close the fire doors nor his failure to water Smith's plants were causes of the fire or of the plants' death. The relationist view, of course, gets these cases right: Brown's omissions are not causes, because *no* absences—and therefore no omissions—are. But nonrelationism gets these negative causal judgments wrong. According to the relationist, the commonsense claim that Jones's omission caused the fire is false; according to the nonrelationist, the commonsense claim that *Brown's* omission did *not* cause the fire is false.

So when it comes to doing justice to commonsense intuitions about causation by absence, neither view fares very well. Luckily, I don't think this matters very much. Briefly, the reason it doesn't much matter is this. As I'll argue in sections 2 and 3, commonsense intuitions about which absences are causes and which aren't are highly dependent on judgments that it would be highly implausible to suppose correspond to any real worldly difference at the level of the metaphysics of causation. For instance, sometimes common sense judges the *moral* status of an absence to be relevant to its causal status. But no philosopher working within the tradition I'm concerned with here thinks that the *truth* conditions for causal claims contain a moral element. It follows that whatever we think about whether or not causation is a relation, we're going to have to concede that common sense is just wrong when it takes, say, moral differences to determine causal differences. There is no genuine causal difference between those cases that common sense judges to be cases of causation by absence and those that it judges *not* to be cases of causation by absence. Hence both the relationist and the nonrelationist must agree, and with good reason, that commonsense judgments about causation by absence are often mistaken—though of course they disagree about *which* commonsense judgments are mistaken.

In sections 4 and 5, I show how absences can figure in causal explanations even though they do not cause anything. So although it is false to say that Smith's failure to water her office plants caused their death, it is nevertheless true to say that the plants died because Smith failed to water them. Here I appeal to Lewis's analysis of causal explanation, according to which to explain an event is to provide information about its causal history: information that need not be restricted merely to citing the event's causes. I claim that common sense judges some absences to be causes because it fails to distinguish between causation and causal explanation, and (in section 6)

that the sorts of distinctions discussed in sections 2 and 3 that ground commonsense assertions and denials of causation by absence are best seen as distinctions between explanatorily salient and nonsalient absences. The moral is that commonsense intuitions about causation by absence are no more damaging to the relationist than they are to the nonrelationist; hence, those intuitions provide no good reason for abandoning the network model of causation.

2 Causation by Absence and Common Sense

In what follows, I'm going to assume that absences that are causes—if there are any—are at least necessary conditions of their effects. I dare say this is a false assumption: Maybe (again, if there is any causation by absence at all) there are cases where, had the absent event occurred, the effect in question would still have had some chance of coming about. I ignore this possibility for the sake of simplicity.

With the assumption that absences are at least necessary conditions of what they cause in place, we might try the following definition of causation by absence:

(I) The nonoccurrence of an event of type *A* caused event *b* if and only if, had an *A*-type event occurred, *b* would not have occurred.

This definition is different from both Lewis's and Dowe's,⁶ and may well be inadequate for reasons other than the one discussed below. However, since both Lewis's and Dowe's definitions can easily be seen to fall prey to the objection too, it will make things easier if we make do with the simpler definition (I).

Not surprisingly, (I) gets the right answer in cases that common sense judges to be genuine cases of causation by absence. Here are some examples.⁷ Flora normally gratuitously waters her neighbor's orchids. But she stops watering them, and they die. Common sense judges that Flora's omission was a cause of the orchids' death, since had she watered them, as she usually does, the orchids would not have died. Second example: *Z*'s dog is bitten by an insect, contracting an eye disease as a result, which *Z* ignores. The dog loses its sight. Intuitively, *Z*'s negligence caused the dog's blindness; and again, the definition gets this right: Had *Z* not ignored the eye disease, the dog would not have lost its sight. Third example: An old tree decays and falls to the ground, crushing some rare wild flowers. The park ranger failed to inspect the tree, thereby making himself a cause of the flowers' death. Again, (I) gets this right: Had the park ranger inspected the tree (and thus had it carefully removed rather than allowing it to fall), the flowers would have survived. Fourth example: A geranium survives the London winter because there is no frost. Intuitively the absence of frost was a cause of its survival; and, once again, this judgment is in accordance with (I).

So far so good. The trouble is that our definition is too inclusive: It renders all sorts of absences as causes that common sense does not *recognize* as causes. Commonsense intuition identifies Flora's failure to water the orchids as a cause of their death, but not the failure of the other neighbors, and certainly not the failure of people on the other side of the world who neither know nor care about the orchids—even though it's perfectly true of each of them that had they watered the orchids, they would not have died. Similarly, while common sense judges the dog owner's or the park ranger's omission to be a cause of the dog's blindness or the crushing of the flowers, it does not judge the omissions of others as causes of those events, even though the relevant counterfactuals are true too. And similarly for the geranium example: We would ordinarily judge that the absence of a hungry geranium-eating koala was no part of the causal history of the geranium's survival, even though had such a koala been present the geranium would not have survived.

Why is this? What grounds common sense's discrimination between absences that are and are not causes of an event? Hart and Honoré, in their book *Causation in the Law*—a classic text for those interested in commonsense causal judgments—claim that what makes us single out one omission as a cause but not another is *abnormality*. We regard Flora's failure to water her neighbor's orchids as a cause of their death because Flora's failure is abnormal: She normally *does* water them. On the other hand, there is nothing abnormal in the failure of the other neighbors, or of other people in other countries, to water the orchids; hence, those failures are not causes of the orchids' death.⁸

The abnormality criterion seems to work for the orchid case, and it seems plausible to suppose that it also works for cases like the geranium case: cases that are not omissions by human agents. The absence of frost in winter in London is quite unusual, and hence according to Hart and Honoré's criterion may count as a cause of the geranium's survival. The absence of hungry geranium-eating koalas, on the other hand, is perfectly normal, and hence does not qualify as a cause of the geranium's survival.

Stapleton claims, however—and I think she's right—that there are many cases in which the abnormality criterion fails to explain our commonsense causal judgments. In the case of *Z*'s dog, for instance, she says that we would regard *Z*'s conduct as a cause “because against the backdrop of *Z*'s ownership of the dog we expect *Z* to have acted, not because of expectations generated by past conduct but for moral reasons” (1994, p. 122). And in the park ranger case, she says, we single out the park ranger and nobody else because she, and nobody else, has a legal duty to inspect the tree. The point here is that we take these causal judgments to be correct even if the relevant omissions are not in any way abnormal: The judgments stand even if *Z* is

generally very bad at looking after his dog, or if it is the park ranger's first day on the job.

Stapleton also claims that whether an omission is to count as a cause can depend on epistemic features. Suppose, for instance, that a certain drug in fact has harmful side effects, but that this risk is unforeseeable: We could not reasonably expect the drug manufacturer to have known about it. In such a case, she says, we should not say "it was the defendant-manufacturer's failure to warn of the unforeseeable risk which caused the injury—clearly a silly idea which no common sense version of causation could accommodate" (ibid., p. 124). On the other hand, if the manufacturer *did* know of the risk, or if we could reasonably have expected them to have found out about the risk, then we *would* say that their failure to warn consumers of the risk was a cause of the side effects.

So it seems that common sense singles out an absence as a cause when, and only when, it stands out in some way—either from what normally happens, or from some norm the absence of which (generally an omission) counts as a violation. The norm might be moral, legal, or epistemic, as Stapleton's examples illustrate; but other sorts of norm may well play a similar role. Owen's failure to get to the ball after Beckham's cross into the penalty area counts as a cause of England's defeat, but Seaman's failure to do so does not—even though, had Seaman (the goalkeeper) somehow managed to get to the ball, he would undoubtedly have scored and England would have won. Why? Because Owen is an attacker and Seaman is the goalkeeper; it's an attacker's job, and not the goalkeeper's, to score goals from crosses into the penalty area.

What all this points toward is a definition of causation by absence that goes something like this:

- (II) The absence of an *A*-type event caused *b* if and only if
- (i) *b* counterfactually depends on the absence: Had an *A*-type event occurred, *b* would not have occurred; and
 - (ii) the absence of an *A*-type event is *either* abnormal *or* violates some moral, legal, epistemic, or other norm.

This definition, I think, does justice to as many commonsense causal judgments about absences as anyone could wish for. The trouble is that although it works *qua* linguistic analysis of the ordinary *concept* of causation by absence, it doesn't look like the sort of analysis we ought to be giving of the *metaphysics* of causation by absence.

Take the violation-of-norms part of the definition. If we take the definition to give the *truth conditions* of causation by absence claims, then causal facts about absences

depend in part on normative facts: facts about whether a moral or epistemic or other norm has been violated. But nobody within the tradition of the metaphysics of causation that I'm concerned with here thinks that causal facts depend on human-dependent norms.

Even if taking causal facts to depend on normative facts weren't in itself such a bad thing, it would in any case make the truth of causal claims turn out to be a relative matter. For instance, you and I might differ in our epistemic standards, so that I count the side effects of a drug as foreseeable and you don't; hence I count the drug manufacturer's failure to warn consumers as a cause of their side effects but you do not. Either there is some absolute epistemic standard that marks out what someone can reasonably be expected to foresee—which seems wildly implausible—or the causal status of the failure is relative to different standards.

A similar point applies to the abnormality part of the definition. For one thing, how often something has to happen in order to count as "normal"—and hence for its absence to count as abnormal—is always going to be rather fluid. And for another, the same absence might count as normal relative to one kind of regularity and abnormal relative to another. Suppose, for instance, that Flora is in the habit of doing gratuitous repetitive favors for her neighbors, like watering their gardens or washing their cars, but that she invariably stops doing them after a month or so. Now is Flora's failure to water her neighbor's orchids after thirty days of watering them normal or abnormal? There doesn't seem to be any principled way of answering that question, so it looks as though the causal status of her omission is simply indeterminate.

It follows that if we want facts about causation to be reasonably determinate and not relative to extraneous facts about whether a putative cause happens to count as abnormal, immoral, illegal, or whatever, then any account of the metaphysics of causation by absence is going to have to be pretty revisionary: No adequate account of what in the world makes causal claims about absences true or false is going to be able to condone all, or even most, of the verdicts given by commonsense usage.

3 A Way Out for the Nonrelationist?

The definition (II) of causation by absence I proposed in the previous section had two conditions on when the absence of an *A*-type event is to count as a cause of *b*: the counterfactual condition—had an *A*-type event occurred, *b* would not have occurred—and the condition that the absence be either abnormal or the violation of some norm. I claimed that no respectable theory of the truth conditions of causal claims ought to respect the second of these conditions; which leaves us with the first,

counterfactual condition, that is, with (I). As we've already seen, an analysis of causation by absence that consisted solely in this condition would be far too inclusive. It would render all of us causally responsible for indefinitely many events happening right now, all over the world—and doubtless beyond. One of the causes of Maite's drinking her coffee in Mexico City just now was my failure to shoot her before she had a chance to put the kettle on. One of the causes of your reading these words right now is the absence of a lion from the room. And so on.

But, in principle at least, the nonrelationist who thinks there really is such a thing as causation by absence is not required to think that there's *that* much causation by absence, for she might be able to invent a more restrictive definition that will do a lot better than what we've got so far. That is to say, maybe our original definition (I) can be supplemented by some other clause that rules out the spurious cases of causation by absence like the one I just gave, without ruling out the allegedly nonspurious cases we want to keep, but which doesn't appeal to normative features of absences.

Well, on behalf of my nonrelationist opponent, I generously offer the following definition. It's the best I can do—but not, as we'll see, good enough:

(III) The absence of an *A*-type event caused *b* if and only if:

- (i) if an *A*-type event *had* occurred, *b* would not have occurred; and
- (ii) an *A*-type event occurs at a world that is *reasonably close* to the actual world.

The definition discounts vast numbers of absences from being causes on the grounds that worlds where the absent event occurs are very distant worlds. For instance, my failure to shoot Maite no longer counts as a cause of her drinking her coffee because, I'm happy to say, a world where I *do* shoot her is a very distant world indeed.

One might even go so far as to claim that the definition gets *all* the cases right (though I'll argue in a moment that it doesn't). We might try to claim, for instance, that what makes Flora's failure to water her neighbor's orchids a cause of their death but not the failure of her other neighbors is not, as we first thought, the fact that Flora's failure, unlike those of the other neighbors, is abnormal; rather, the difference is that Flora is more strongly disposed to water the orchids than are her neighbors—as evidenced by the differences in her and their past behavior. So a world where Flora waters the orchids is closer to our own than is a world when another neighbor does so. *Z*'s failure to attend to his dog's eye infection, as opposed to our failure, counts as a cause of its blindness because *Z* was in a position to do something about it and we were not: Hence a world where *Z* takes his dog to the vet is reasonably close to our own, and one where one of us takes the dog to the vet is not.

The thought here, then, is that abnormality and violation of norms are not part of the *truth* conditions of causation by absence claims. Rather, those considerations inform our judgments about how reasonable it is—that is to say, epistemically reasonable—to expect the absent event to happen; and thus how reasonable it is to suppose that the event in question happens at a reasonably nearby world. For example, we generally assume that dog owners are disposed to behave in a way conducive to the welfare of their pets, whereas people in general are not particularly disposed to go out of their way to behave in a way conducive to the welfare of other people's pets. And we can construe this assumption as an assumption about closeness of worlds: If *Z* is strongly disposed to keep a close eye on his dog's welfare and I am not, then a world where *Z* manifests this strong disposition and takes the dog to the vet is plausibly closer than a world where I, who have no disposition whatever to do so, take *Z*'s dog to the vet for him. Hence *Z*'s omission is a cause of the dog's blindness but mine is not.

Things seem to be looking up for the nonrelationist; but unfortunately I don't think the situation is as rosy as I've made it look. I have two objections to the suggested definition. First, I think that commonsense intuition would still discriminate between, say, *Z*'s omission and mine even if we knew perfectly well that *Z* is a terrible pet owner and consistently fails to look after his dog properly. Even if we think that *Z* has no disposition whatever to take the dog to the vet, and hence that a world where he does so is just as distant as a world where someone else does it for him, we still think of *Z*'s omission and nobody else's as a cause of the dog's blindness.

Similarly, we still judge a drug company's failure to warn consumers of foreseeable side effects as a cause of their illness even if we know that the company is extremely disreputable and rarely bothers to carry out the proper tests. I take it that commonsense simply doesn't endorse the view that if you're negligent enough—if your disposition to behave in accordance with norms is weak enough—your negligence literally won't have any effects.

These observations support my initial claim that the norms are doing the work by themselves, as it were, rather than merely informing our judgments about closeness of worlds. Still, I'm basing this conclusion on some disputable claims about commonsense intuition that you may not share, so this isn't a decisive objection.

The second objection, I think, is more telling: I can see no sensible way of specifying what "reasonably close" amounts to. It's one thing to judge *relative* similarity of worlds—to judge that, say, a world where Flora waters her neighbor's orchids is closer than any world where someone else waters them. But it's quite another thing to try and impose a *metric* on this ordering, so that any *A*-world within a certain distance from the actual world is close enough for the absence of *A* to count as a

cause, whereas any *A*-world outside that distance is too far away for the absence of *A* to count as a cause.

Even if there *were* a way of specifying what “reasonably close” amounts to, there are two further problems. One is that any choice of maximum distance is going to be entirely arbitrary; and the question of which absences are causes and which are not ought not to be decided by a piece of arbitrary stipulation. The other problem is that whereas relative similarity of worlds might explain why one absence and not another counts as a cause in any particular case—*Z*’s taking his dog to the vet happens in a close enough world, say, but not my taking it—I see no reason to suppose that the same standards will apply across all cases. What counts as close enough for causation in one case might not count as close enough for causation in another.

All of these problems for the proposed definition (III), I think, have the same source: There just *isn’t* any objective feature that some absences have and others lack in virtue of which some absences are causes and others are not. So *any* definition of causation by absence that seeks to provide a principled distinction between absences that are and are not causes is bound to fail: No such definition will succeed in carving nature at its joints. If this is right, then the nonrelationist is going to have to concede that there just *is* no principled reason to regard Flora’s failure to water the orchids as a cause of their death but not to regard my failure to shoot Maite as a cause of her coffee-drinking. And this, I suggest, hardly puts nonrelationism ahead of relationism when it comes to doing justice to commonsense intuitions.

In fact, there is another possible way out of the problem for the nonrelationist, and that is to deny that the alleged distinction in our commonsense talk and thought between absences that are and are not causes really exists. I have been claiming that, according to commonsense intuition, Flora’s failure to water her neighbor’s orchids was a cause of their death, but the failure of other people—people who neither know nor care about the orchids—was not. But one might object that this is not really the position endorsed by common sense. Rather, in ordinary circumstances we fail to identify the failures of other people as causes of the orchids’ death, in the sense that it does not occur to us to mention those failures’ part in the causal history of the death of the orchids (since in ordinary conversational circumstances it would be inappropriate for us to do so). But to fail to mention that those other failures are causes is not to hold that they are *not* causes.

No doubt there is some truth in this suggestion. Perhaps, if asked directly whether the failures of other people to water the orchids was a cause of their death, at least some people would say yes—presumably because they would appreciate the similarity between Flora’s failure, which they do count as a cause, and the failures of those other people.

I see no reason to suppose, however, that the point generalizes. The number of possible events, or combinations of events, that are such that, had they occurred, the orchids would not have died, is absolutely enormous. The plants would have survived, for example, if Flora's neighbor had installed a sprinkler system that then got activated accidentally while she was away, or if her roof had started leaking at a point just above the orchids during a rainstorm, or if a cow had somehow entered her house and flicked a nearby glass of water onto the orchids with its tail, or—. I do not think that most people would happily accept that the failure of each of these events to occur was equally a cause of the orchids' death. Of course, this is an empirical claim about what people are ordinarily inclined to judge. But I have not come across any evidence to suggest that the claim is false.

4 Causation and Causal Explanation

I shall return to the commonsense distinction between absences that are and are not causes in section 6. In this section and the next, however, I set this issue aside in order to focus on another distinction: the distinction between causation and causal explanation. In this section I defend the view that not all causal explanations are reports of causation: The explanans of a causal explanation need not stand to the explanandum as cause to effect. And in section 5 I defend the view that causal explanations that involve facts about absences can be seen as explanations of just this kind: We do not need absences as causes in order for facts about absences to be the explananda of causal explanations.

In "Causal Relations" (1967), Davidson argues for a distinction between causation and causal explanation by concentrating on the logical form of causal statements. He argues, using a version of the "Slingshot" example, that causation is a relation between events rather than facts—in other words, that the canonical form of causal statements is "*c* caused *e*," where *c* and *e* are events and *caused* is a two-place relation, rather than "*E* because *C*," where *C* and *E* are facts and "because" is a sentential connective. Davidson reserves the "*E* because *C*" locution for causal explanation. For Davidson, then, no causal explanations are themselves causal claims, since they simply do not have the right logical form.

In opposition to Davidson, Mellor has long maintained that *all* causal explanations are in fact causal claims. For Mellor, facts are the most basic kind of causal relata, and the canonical form of causal statements is "*E* because *C*."⁹ According to Mellor, causation *sometimes* also relates events—but only sometimes. When it's true that the match lit because I struck it, it's also true that the striking caused the

lighting; but when it's true that, say, Kim has no children because she used contraception, there is no corresponding true statement of the form "*c* caused *e*."¹⁰

If absences are to figure in causal explanations without doing any causing, there must be a distinction between causation and causal explanation: Some causal explanations cannot be reports of causation. But this latter claim needs to be defended against the following objection: How, one might ask, can a causal explanation be genuinely *causal* if the explanans doesn't stand to the explanandum as cause to effect? Or, as Mellor puts it, "how can facts explain other facts causally without causing them?" (1995, p. 130). Well, here I want to appeal to Lewis's theory of explanation. For Lewis, "to explain an event is to provide some information about its causal history" (1986b, p. 217); and the causal history of an event, he says, is "a relational structure" (*ibid.*, p. 216). Not surprisingly, then, Lewis's account of explanation is tailor-made to fit the network model of causation. For present purposes, the most important feature of Lewis's account of explanation is that it does *not* amount to the view that every explanation involves picking out a cause (or some causes) of an event; the way in which causal facts enter into an explanation can be more complex than that. One can give information about an event's causal history in all sorts of other ways—by saying, for instance, that certain events or kinds of event do *not* figure in its causal history, or by saying that an event of such-and-such *kind* occurred, rather than that some *particular* event occurred. The moral here, then, is that something can be the explanans of a causal explanation without itself being a cause of the event cited in the explanandum.

For example, suppose that Lee Harvey Oswald really did shoot JFK. Then the following three sentences are all true:

- (1) Oswald's shot caused JFK's death.
- (2) JFK died because Oswald shot him.
- (3) JFK died because somebody shot him.

On Mellor's view, all three are *causal* truths. On Lewis's view (as expressed in Lewis 1973a, 1986b,c), only the first is, strictly speaking, a causal truth; the second and third are causal explanations rather than reports of causation. It's pretty obvious what makes (2) a *causal* explanation: Each of the explanans and the explanandum asserts that a particular event (Oswald's shot and JFK's death respectively) occurred, and those two events are in fact causally related. The case of (3) is a little more complicated: Here, the explanans does not assert that a *particular* event occurred, for there is no event essentially describable as *someone's* shooting JFK. Such an event would be disjunctive: It would be the event of Lee Harvey Oswald's shooting or the

man on the grassy knoll's shooting or Jackie Kennedy's shooting or ... and so on. And there is no more such an event than there is an event of my-birthday-party-or-your-morning-bath.¹¹ Rather, the explanans of (3) asserts that there was some event or other that was a shooting of JFK by someone. Hence the explanans of (3) does not stand to the explanandum as cause to effect. Nonetheless, (3) still counts as a true causal explanation because, although it doesn't tell us *which* event caused JFK's death, it tells us something about the death's causal history, namely, that it included a shooting by someone.

The notion of providing information about an event's causal history is further expanded by Jackson and Pettit's account of program and process explanations.¹² This account is designed to show how multiply realizable properties (like functional and dispositional properties) can figure in true causal explanations without being "causally efficacious." For example, suppose that to be fragile is to be disposed to break when dropped, and that a glass is dropped and duly breaks. We recognize that it is the glass's molecular structure that causes it to break, and not its fragility; but we do not want to conclude that the fragility cannot figure in a worthwhile *explanation* of why it broke. (Likewise, we do not want the view that mental states are functional states to preclude the possibility of explaining, say, actions in terms of beliefs and desires.)

Jackson and Pettit's solution to the problem is to distinguish between "process" and "program" explanation. "Process" explanation is explanation in terms of actual underlying physical processes: an explanation of why the glass broke in terms of molecular structure, say, or an explanation of why I went to the shop in terms of my neural processes. Program explanations, on the other hand, tell us not what the actual underlying processes were, but rather that those processes satisfy a particular functional or dispositional or higher-order description. The presence of fragility "programs for" the presence of a causally efficacious molecular property. Although it is the actually present causally efficacious property, and not the programming property, that figures in the causal history of the breakage, the program explanation still tells us something *about* that causal history—namely, that it involves *some efficacious property or other* for which fragility programs.

As Jackson and Pettit say:

The process story tells us about how the history actually went: say that such and such particular decaying atoms were responsible for the radiation. A program account tells us about how that history might have been ... telling us for example that in any relevantly similar situation, as in the original situation itself, the fact that some atoms are decaying means that there will be a property realized—that involving the decay of such and such particular atoms—which is sufficient in the circumstances to produce radiation. In the actual world it was this, that and

the other atom which decayed and led to the radiation but in possible worlds where their place is taken by other atoms, the radiation still occurs.¹³

If we adopt Lewis's analysis of causal explanation, then, we can distinguish between causation and causal explanation and still be able to say why causal explanations count as *causal* explanations; for such explanations can give information about the causal history of the event to be explained even though the explanans does not stand to the explanandum as cause to effect.

5 How Can Causally Inert Absences Explain Anything?

In the last section, I defended the view that facts can causally explain without being *causes* of what they explain. In this section, I argue that facts about absences, omissions, and failures do just that. The "because" locution is—or at least sometimes is—an explanatory locution; moreover, as we have seen, "because" claims can be true—can reveal information about causal history—without the explanans standing to the explanandum as cause to effect. So we can repudiate the claim that absences can be causes and perfectly well grant that there are true causal explanations whose explananda concern absences.

First, though, a point about the logical form of causation by absence claims. Much of the time in our everyday talk, we speak as if absences, omissions, and failures are *things*. We say that the void is deadly just as we would say that the Chrysler Building is very tall: The sentence has the same subject-predicate form, and we count it as true, even though there could be no *object* in the world picked out by the definite description "the void." Similarly, we say that Flora's failure to water the orchids caused their death just as we might, in different circumstances, say that Flora's throwing them on the fire, or cutting them up, caused their death. Our everyday causation-by-absence claims often have the logical form of a relational sentence—and we count them as true—even though one of the singular terms flanking the relation does not refer to anything. Absences, omissions, and failures get assimilated to the familiar ontological category of events even though they are *not* events.

So what? Well, suppose for a moment that there really is causation by absence, as the nonrelationist claims. Even so, one ought to think that expressing causation by absence facts using the "*c* caused *e*" locution is at best highly misleading. For such sentences have a relational form, yet whatever it is that makes them true, it is not the obtaining of any relation. The discussion in section 4 suggests an obvious paraphrase of causation-by-absence claims into the "because" locution: A more ontologically perspicuous way of saying that Flora's failure to water the orchids caused their death

is to say that the orchids died—or that their death occurred—because Flora failed to water them.

The relationist, on the other hand, cannot claim that “the orchids died because Flora failed to water them” is a paraphrase of “Flora’s failure caused the orchids’ death,” for, although we can perhaps paraphrase away the relational form of the latter sentence, we cannot paraphrase away the fact that it is a *causal* claim; and according to the relationist, absences do not *cause* anything. So if we want—as I want—to hold that the sentence “the orchids died because Flora failed to water them” is explanatory but *not* causal, it cannot be a paraphrase of the relational sentence.

Still, I agree with the nonrelationist (as I have portrayed her) on this much: We would do better to use the “because” locution than the “caused” locution when we are talking about absences. The nonrelationist should think so because the “because” locution is less misleading, whereas I think so because the “because” locution will allow us to say something true, whereas the “caused” locution will not.

I say that common sense is just mistaken when it asserts that an absence or an omission *caused* some event. It’s not an especially bad mistake. Often we move between the “*E* because *C*” and “*c* caused *e*” locutions without going wrong: It doesn’t much matter whether I say “the match lit because I struck it” or instead “my striking the match caused it to light”; or whether I say “the crash occurred because there was an avalanche” or instead “the avalanche caused the crash.” Often causal explanations go hand in hand with causal relations between events. Often, but not always. When I say “Flora’s failure to water the plants caused their death” instead of “the death occurred because Flora failed to water them,” I say something false instead of something true. It’s not surprising, nor a matter for particular concern, that we make this error in our everyday talk; no serious harm is done. But I see no reason why, from a philosophical perspective, we should not rule the move out of order.

It remains to be shown, of course, that “because” claims involving absences really can give information about the causal history of the event to be explained without the absence being a *cause* of that event. What sort of information does such an explanation give us about the event whose occurrence we want to explain? Well, when I say that the orchids’ death occurred because Flora failed to water them, you learn something minimal about the death’s causal history: that it did not include an event of Flora’s watering the orchids. But you also learn something about the causal structure of nearby worlds where Flora *didn’t* fail to water the orchids—namely, that the causal processes that ensued at those worlds did not cause (or perhaps might not have caused) the orchids’ death. And if you doubted the veracity of my explanation, I could fill in more details about how the causal history of the world would have

gone had Flora watered the plants: Water would have been taken up by the roots, sustained the cell structure, and so on. None of this, of course, is information about what causal processes there were in the *actual* world; it is information about what causal processes there are in the closest world(s) where the actually absent event occurs. As with program explanation, the information provided about causal history is *modal* information; and it is information that is explanatorily relevant to the orchids' death.

Similarly for Lewis's void. Suppose that as punishment for her negligence, Flora's neighbor casts her into the deadly void. Flora's blood boils, the air is sucked from her lungs, and so on. If, as I claim, there is no causation by absence, then the void causes none of these unfortunate events. Strictly speaking, the void is not deadly—if deadliness is the capacity to cause death. But Flora's death is not thereby rendered uncaused: There are plenty of positive events going on in her body that *do* cause her death. Nor, on the account offered here, is the void explanatorily irrelevant to Flora's death. It's perfectly true to say that Flora's blood boiled because there were no forces present—forces that ordinarily keep us alive—to stop it. When we cite the void in our explanation of Flora's death, we describe how Flora's causal history *would* have gone had she *not* been cast into the void. We do not say what *actually* caused her death; rather we point out that the sorts of events that would have caused her to remain alive did not occur.

I see no reason, then, why explanations invoking facts about absences should not be seen as genuine causal explanations, even though absences do not cause anything. Of course, I have not shown that this relationist account of the role of absences in causal explanations is any better than a nonrelationist account like Mellor's, according to which facts about absences causally explain in virtue of their being causes. I merely hope to have shown that the relationist about causation is not committed to denying that absences have a legitimate role to play in our explanatory talk.

6 Commonsense Judgments and the Pragmatics of Explanation

It's time to return to the issues discussed in the first three sections of the essay. There I argued that commonsense assertions and denials of causation by absence depend on considerations—normative features, for example—that have no place in an account of the metaphysics of causation. Since, from the perspective of metaphysics, we can draw no relevant distinction between, for example, Flora's failure to water the plants and your failure to do so, or between *Z*'s failure to attend to his dog's eye infection and *Z*'s neighbor's cousin's best friend's failure to do so, the defender of causation by

absence must concede that there is much more causation by absence about than we ordinarily think there is. The denier of causation by absence, on the other hand, must concede that there is much *less* causation by absence than we ordinarily think there is: There is none at all, in fact.

As a denier of causation by absence, what do I say about the commonsense distinction between “Flora’s failure to water the orchids caused their death” (true) and “your failure to water the orchids caused their death” (false)? I say, of course, that both are false. But the corresponding *explanatory* claims (“the orchids died because Flora failed to water them”; “the orchids died because *you* failed to water them”) are both true. However, they do not, in most contexts, count as equally *adequate* explanations. When we explain why the orchids died, we must, if our explanation is to count as adequate as well as true, be sensitive to why the explanation was requested in the first place. In the context where my interlocutor is requesting information that is relevant to the issue of whom to *blame* for the death of the orchids, it would be highly misleading of me to say that they died because you didn’t water them. The *truth* of my utterance doesn’t depend on the moral question of who is to blame; but the *adequacy* of my explanation does, in this context, so depend. Similarly, it’s true to say that I attended the seminar because I wasn’t attacked by a hungry polar bear; it’s just *very* hard to imagine a context within which someone who asked me why I attended would be satisfied with that explanation.

Like the distinction between events and absences, the distinction between a true explanation—a true “because” statement—and an adequate explanation is one that common sense has a tendency to ignore. If you judge that the orchids died because Flora didn’t water them but not because you didn’t water them yourself, or that I attended the seminar because I didn’t have anything better to do and *not* because I was not attacked by a hungry polar bear, you mistake lack of explanatory salience for falsity.

Of course, the believer in causation by absence can tell the same story. I do not claim that my account is any more plausible than that of someone who thinks there *is* causation by absence; I merely hope to have persuaded you that it is no *less* plausible, and hence that it is no objection to the network model of causation that it entails that there is no causation by absence.

The causal history of the world is a mass of causal processes: Events linked by a vast and complex web of causal relations. In order that the causal history of the world should look the way it does look, rather than some other way, there must have been no *extra* events impinging on it—for those extra events would have had effects that would have changed the causal history of the world in various ways. If Godzilla had impinged upon the causal history of the world, that causal history would have

gone very differently. We might even, if circumstances demanded it, want to explain happenings in the world by citing Godzilla's absence (though it's hard to imagine that we should ever want to do so). But I see no need to think of Godzilla's lack of impingement as a kind of causation.

Notes

1. See Steward (1997), chapter 7.
2. See for instance Lewis (1973a, 1986b), and Davidson (1967).
3. Lewis, chapter 10 in this volume, p. 281.
4. See Mellor (1987) and (1995, chapter 11). Steward (1997) also believes that the network model should be abandoned.
5. See Lewis, chapter 10 in this volume, pp. 284–285. Another proposal along similar lines can be found in Dowe (1999).
6. See Lewis, chapter 10 in this volume, p. 284 and Dowe (1999).
7. The first three examples are taken from Stapleton (1994), chapter 6, section 2.
8. See Hart and Honoré (1985), p. 38.
9. Of course, “relata” isn't really the right word, since Mellor agrees with Davidson that “because” does not pick out a causal relation.
10. See Mellor (1995), pp. 162–165.
11. See Lewis (1986d), p. 267.
12. See, for example, Jackson and Pettit (1990).
13. Jackson and Pettit (1990), p. 117.