

On Causing Something To Happen In A Certain Way Without Causing It To Happen*

Carolina Sartorio, University of Wisconsin-Madison

We tend to identify events by when, where, and how they occur, i.e. by the time at which they occur, the spatial location at which they occur, and the manner in which they occur. But, in most cases, we also tend to conceive of the *specific* time, spatial location, and manner in which they occur as only *accidental* to them. For instance, we think that an explosion that occurred at noon could have occurred one minute later, an inch closer to our home, and it could have been a darker shade of blue. Our meeting today could have occurred an hour later, in the room next door, and it could have involved a few more people than it did.¹

Given that many of the properties events have are accidental to them, it is possible to cause an event to have a property without causing the event itself. For instance, I can cause an explosion to be blue without causing the explosion, and I can cause a meeting to be boring without causing the meeting.² Other times, the causing of an event to have an accidental property comes hand in hand with the causing of the event itself. Suppose, for instance, that Assassin shoots Victim, and Victim dies as a result. The particular way in which the death occurred involves a certain amount of blood spilling. Imagine that Victim's death could still have occurred if more or less blood had been spilled, that is, imagine that the particular amount of blood being

* I am specially grateful to Karen Bennett, Juan Comesaña, and Laurie Paul for their helpful comments and criticisms. Thanks also to Adam Elga, Liz Harman, John Hawthorne, Ned Markosian, Ceil Maslen, Daniel Nolan, Josh Parsons, Jonathan Schaffer, Gabriel Uzquiano, Stephen Yablo, the audience at the 2004 Bellingham Summer Philosophy Conference, and the members of the metaphysics and epistemology reading group at the University of Wisconsin-Madison.

¹ At least, this seems to be the right picture of events as causes and effects, which is the topic that concerns me (see Lewis (1986b)). If we make events' essences too specific or "fragile," it is hard to avoid the result that any minute influence in the way an event happened was a cause of the event's happening.

² For the distinction between causing an event and causing an event's "aspect," see Paul (2000). It is possible to capture a similar distinction without stepping outside of the events framework. For instance, instead of saying that one causes the explosion's *blueness* without causing the explosion, one could say that one causes an event that is essentially blue but not an event that is essentially an explosion. According to this view, there are coincident events that differ only in their modal properties. Also, one could put the point in terms of fact-causation. For instance, one could say: it is possible to cause the fact that an explosion was blue without causing the fact that the explosion happened. I wish to remain neutral among these different frameworks.

spilled is only accidental to the death. Then Assassin caused the death to have a property, which is only accidental to it, but he also caused the death *simpliciter*.

When is causing an event to have certain properties also causing it to happen, and when is it not causing it to happen? This is the topic of this paper. In what follows I argue for a principle that gives sufficient conditions under which causing an event to have certain properties is *not* causing the event. The principle has the following features. First, it unifies cases of different types: cases where something is a cause of either *when*, or *where*, or *how* an event occurs without causing the event. Second, the common element is the fact that in all those cases the event in question occurs, intuitively, *not because* of what caused it to have the accidental property, but *in spite of* it. In other words, the principle captures the concept of what I shall call an “inhibitor.” Finally, the principle has the consequence that causation is not a transitive relation. As such, it works as an argument against the transitivity of causation.

Let me stress that the aim of this paper is only to give sufficient conditions for the phenomenon of causing an event to have a property without causing the event, not sufficient and necessary conditions. As we will see, this less ambitious project is already quite challenging, as well as quite fruitful.³

I.

I will start by demarcating the properties of events that I will be concerned with. In a nutshell, I will be concerned with *ordinary* properties of events. As hinted at before, I take ordinary properties of events to be those properties that determine when, where, and how events occur (i.e.

³ One might think that there is an easy way of offering sufficient and necessary conditions, namely this: x causes an event to have an accidental property without causing the event if and only if x doesn't cause any *essential* properties of the event. Now, even if this principle were true, it is not very illuminating. Was Assassin causally responsible for Victim's death? Maybe he was just in case he was causally responsible for some essential property of the death, e.g., its involving Victim's corpse. But, was he causally responsible for that essential property? It seems that the issue of whether he was causally responsible for that property is basically the same as the issue of whether he was causally responsible for the death. If so, the suggested principle is trivial and uninteresting.

their spatiotemporal location and their manner of occurrence), and I will only be concerned with properties of this kind that events have accidentally (i.e. properties that they have but that they could have failed to have).

In most cases, and within certain limits, it is clear what the spatiotemporal location of an event is. But, what is the “manner” of occurrence of an event? Following a Lewis-style approach to events, I will assume that this is determined by what happens in the spatiotemporal region where the event occurs (Lewis (1986b)). This is to say, once we have identified the spatiotemporal region where the event occurs, the way in which the event occurs or its manner of occurrence is determined by what happens within that region. This is intended to rule out *extrinsic* properties that events might have. Extrinsic properties are determined, at least in part, by what happens in other regions (for instance, being the largest explosion of all times is only extrinsic to, say, the Big Bang because it depends on the size of other past, present, and future explosions). Intuitively, extrinsic properties are properties that events have, not purely in virtue of what *they* are like, but at least partly in virtue of what *other things* are like. Intrinsic properties, by contrast, are properties that events have just in virtue of what they are like.

Another restriction concerns “positive” versus “negative” properties. Should we say, e.g., of a blue explosion that it is also not red, not green, etc.? Intuitively, a blue explosion only *lacks* the property of being red, or being green: it doesn’t *possess*, in addition, another property of being not-red, or not-green. At any rate, to the extent that “negative” properties are genuine properties, I want to set them aside. For, again, intuitively, negative properties don’t determine what an event is like; rather, they determine what the event *isn’t* like.⁴

In sum, I will be concerned with accidental, ordinary properties of events. Ordinary properties of events include spatiotemporal location as well as intrinsic, non-negative properties. To distinguish them from other properties events might have, I will call them *features*.

⁴ Another plausible restriction is to properties that aren’t too “disjunctive”: again, widely disjunctive properties are not ordinary properties of events.

II.

I now turn to discussing under what conditions causing an event to have certain features is not causing the event. Let's start with a particular feature: time. Imagine the following scenario:

Delaying Samaritan: Sam notices that Assassin is about to shoot Victim and attempts to disable him. She fails; however, as a result of her intervention, Assassin's shooting and Victim's death are delayed by one second. (The rest of the details of the death are the same.)

Sam is a delayer with respect to the death: she makes Victim's death happen one second later than it would have happened otherwise. Thus Sam is causally responsible for a feature of the death: a temporal feature. But, intuitively, Sam doesn't cause the death: intuitively, the death doesn't happen *because* of Sam's intervention but *despite* Sam's intervention.⁵ (By contrast, hasteners are such that we typically *do* regard them as causes: if I make a death happen earlier rather than later, there is a temptation to say that the death occurred because of what I did—certainly not despite what I did—and thus there is a temptation to say that I cause the death. More on hasteners and the apparent hasteners/delayers asymmetry later.)

Generally, then, delayers don't strike us as causes because they contribute only “negatively” to the occurrence of the effect, in some loose sense that would have to be spelled out.⁶ This is the case with Delaying Samaritan: in Delaying Samaritan, Sam only contributes “negatively” to the occurrence of Victim's death and thus doesn't seem to be a cause.

⁵ This claim is likely to rely on the assumption that Victim's actual death would still have occurred if Sam hadn't intervened. This is how I will understand delayers: as I will understand them, delayers are such that, if they hadn't occurred, the *same* event that actually occurs would still have occurred, but it would have occurred at an earlier time.

⁶ There are some delayers for which this isn't the case. Imagine that Assassin1 shoots a bullet at Victim. A second later, Assassin2 shoots a second bullet that deflects Assassin1's bullet away from its path and

But, how can we make this idea more precise? I will start by looking at a suggestion by Penelope Mackie. Mackie suggested the following: typically, a delayer doesn't strike us as a cause of an outcome because it causes the outcome to happen at a time (a later time) *only by* causing it *not* to happen at other times (earlier times) (Mackie (1992)). For short, a delayer causes *only by preventing*, and causing only by preventing is not causing. Causing only by preventing is not causing in the sense that causing an event to occur *later* only by preventing it from occurring *earlier* is not causing the event *itself* to occur. In Delaying Samaritan, for instance, Sam causes the death to happen later *only by preventing* it from happening earlier. Thus, she doesn't cause the death. (Mackie claims that, by contrast, the opposite is generally true of hasteners: instead of causing only by preventing, they prevent only by causing. For instance, if I make a death to happen earlier, then I prevent it from happening later only by causing it to happen earlier. More on hasteners later.)

Clearly, the central component of the account is the “only by” locution. Unfortunately, Mackie says that she doesn't know how to account for it and leaves it as a primitive. Now, although I do find the notion intuitive at some basic level, I think that its application to specific situations is probably trickier than Mackie thinks. Hence, I will attempt to say more about how we should understand it.⁷

I will argue that we should understand the “only by” locution as expressing a kind of supervenience or noncausal dependence relation. The relation in question obtains, not between ordinary events or facts, but between episodes of causing or causal facts—for short, “causings.” For instance, on this proposal the claim: “Sam caused the death to happen later only by causing it not to happen earlier” should be understood as: “Her causing the death to happen later noncausally depends on her causing the death not to happen earlier.”

then goes on to knock Victim down. Then Assassin2's shooting causes Victim's death, even if it delays it by a second. (See also Mackie's examples of atypical delayers in her (1992).) An advantage of the theory developed here is that it allows for such delayers to be causes (see n. 14).

⁷ I am not sure whether Mackie would find the account satisfying on all counts, but I do think that it gets her paradigm cases right.

Kim has discussed relations of noncausal dependence between causings in his (1974). According to Kim, there is noncausal dependence between causings in the following cases. First, suppose that I open the window by turning the knob. Kim then claims that my causing the window to open noncausally depends on my causing the knob to turn. Second, consider someone who was causally responsible for Socrates' death. Presumably, that person was also causally responsible for Xanthippe's widowhood.⁸ Kim then claims that that person's causing Xanthippe's widowhood noncausally depends on his causing Socrates' death.

In virtue of what do these relations obtain, according to Kim? He claims that they are based on more basic dependence relations between ordinary events or facts. Consider, first, the example of the knob and the window. Here the two ordinary events in question are the knob's turning and the window's opening. Clearly, there is a dependence relation between them. In this case, the dependence is causal: the knob's turning causes the window's opening. Kim claims that this dependence relation at the level of the events grounds another dependence relation at the level of the causings of those events: given that the window's opening depends on the knob's turning, *my causing* the window to open depends on *my causing* the knob to turn.

Now, we have to be careful about how to understand this claim. Kim is clearly thinking of a case where my only contribution to the window opening is my turning the knob; for instance, it is not the case that, at the same time that I turned the knob, I removed an obstacle that was preventing the window from being open. Although other things were required, apart from my turning the knob, for the window to open (e.g., there had to be nothing obstructing the window), none of these extra conditions concern me. In other words, I take Kim's idea to be the following: in the circumstances, the window's opening depended on the knob's turning; in addition, I was

⁸ Someone like Lewis would object to this claim on the grounds that Xanthippe's becoming a widow is not something that can enter in causal relations because it's "too extrinsic" (whether it occurs in a region would depend on what happens in other spatiotemporal regions, namely, the regions containing Socrates); see Lewis (1986b). However, the example still helps to illustrate the concept of noncausal dependence between causings.

only causally responsible for the knob's turning, and not for the fact that any of the circumstances obtained; hence, *my* causing the window to open depended on *my* causing the knob to turn.

Finally, notice that, whereas the dependence between the events of the knob's turning and the window's opening is causal, the dependence between *my causing* the knob to turn and *my causing* the window to open is not causal. For, surely, *my causing* the knob to turn did not itself *cause* *my causing* the window to open.

The case of Socrates and Xanthippe is only slightly different. The sole difference is that here the relation between the more basic events or facts is itself noncausal: Xanthippe's widowhood is not caused by Socrates' death; however, it still depends on it: Xanthippe became a widow in virtue of the fact that Socrates' death occurred (and not vice versa). Again, we should understand the dependence between them as: in the circumstances (i.e. given that Xanthippe was married to Socrates), Xanthippe's widowhood depended on Socrates' death. In turn, this dependence grounds the dependence between the causings. Take someone who was responsible for Socrates' death but not for the fact that Xanthippe married Socrates: given that Xanthippe's widowhood depended on Socrates' death, that person's causing Xanthippe's widowhood depended on his causing Socrates' death. Again, this dependence between the causings is noncausal.

I suggest that we use Kim's concept of noncausal dependence between causings to understand Mackie's "only by" locution. In the example of the knob and the window, Kim's suggestion that there is noncausal dependence between *my causing* the knob to turn and *my causing* the window to open can be restated as: I caused the window to open *only by* causing the knob to turn. In the example of Socrates and Xanthippe, Kim's suggestion that there is noncausal dependence between the causing of Socrates' death and the causing of Xanthippe's widowhood can be restated as: someone caused Xanthippe's widowhood *only by* causing Socrates' death.

Now recall Delaying Samaritan: in Delaying Samaritan, Sam delays the occurrence of the death by just one second without causing the death. Following Mackie, we want to be able to

say that Sam causes the death to happen later *only by* causing it not to happen earlier. What this means, in Kimean terms, is that her causing the death to happen later noncausally depended on her causing it not to happen earlier.

In virtue of what is this true? As with Kim's examples, we should probably start with a more basic dependence relation between ordinary events or facts. Given that Sam delayed the death by tampering with Assassin's actions, I suggest that we focus on these two facts: the fact that Assassin didn't shoot earlier and the fact that he shot later. Clearly, certain things had to happen in order for Assassin to shoot later given that he didn't shoot earlier (for instance, Assassin had to still want to shoot later if he couldn't do so earlier, and there had to be no further obstacles to his doing so later). These are the circumstances in which Assassin's shooting later depended on his not shooting earlier. Now, in this case it might not be totally clear if the dependence between these facts is causal or noncausal: Was Assassin's not shooting earlier a *cause* of his shooting later, in the circumstances, or was it only a precondition for its happening later? Different people would have different opinions on this matter, but we needn't settle this issue; after all, as we have seen, it is enough that there be a *dependence* relation between the two, whether it is causal or noncausal doesn't matter. (Thus, in the example of the knob and the window the dependence between the basic facts was causal, but in the example of Socrates and Xanthippe it wasn't.)

Now, by assumption, Sam was not causally responsible for the circumstances in which Assassin's shooting later depended on his not shooting earlier (for instance, she wasn't responsible for Assassin's tenacity). Hence, *Sam's* only contribution to Assassin's shooting later was her causing him not to shoot earlier. Thus, in Kimean terms, Sam's causing Assassin to shoot later noncausally depends on her causing him not to shoot earlier. Or, in Mackie's terms, Sam caused Assassin to shoot later *only by* preventing him from shooting earlier. Presumably,

then, Sam caused the death to happen later only by preventing it from happening earlier. This is the result we wanted.^{9,10}

The Kimean account of the “only by” locution is the first way in which Mackie’s account can be improved. The second way in which it can be improved is by generalizing it to other features of events. I turn to this in the next section.

III.

Mackie’s idea was that delayers don’t strike us as causes because they typically cause only by preventing: they cause an event to happen later only by preventing it from happening earlier. That is, delayers cause only by preventing as far as *time* is concerned. Now let us try to generalize this idea to all features of events (i.e., besides temporal location, spatial location and manner of occurrence). Roughly, whereas the principle for time says that something is not a cause of an event if it causes only by preventing as far as the time of occurrence of the event is concerned, the generalized version of the principle says that something is not a cause of an event if it causes only by preventing as far as *any* feature of the event is concerned. More precisely, the generalized version of the principle reads:

Causing by Preventing: If, for every feature F that X causes an event E to have, there is another feature G such that X causes E to have F *only by* causing it *not* to have G, then X is not a cause of E.

⁹ Strictly speaking, the parallel with Kim’s examples only takes us to the conclusion that Sam caused Assassin to shoot later only by preventing him from shooting earlier. However, it is natural to conclude from there that Sam also caused *the death* to happen later only by preventing it from happening earlier. For, if Sam’s only contribution to Assassin’s shooting later was her preventing him from shooting earlier, and if this accounted for the death’s happening later rather than earlier, then it seems that Sam’s only contribution to the death’s happening later was her preventing it from happening earlier.

¹⁰ With typical hasteners, the opposite is generally true. Suppose that I make a death happen earlier rather than later. Then it seems that the death’s not happening later depends on its happening earlier (since, if it happened earlier, it cannot happen again). Thus, I prevent the later occurrence of the death only by causing its earlier occurrence. That is, whereas typical delayers cause only by preventing, typical hasteners prevent only by causing.

In other words, something is not a cause of an event when it causes the event to have certain features only secondarily, merely by preventing other features.

To see how the principle applies to specific cases, imagine the following scenario:

Disabling Samaritan: A criminal organization has two assassins lined up for each crime: if Assassin1 fails, then Assassin2 steps in and does the job. On a certain occasion where Assassin1 is about to shoot Victim, Sam intervenes and disables him. Assassin2 then steps in and kills Victim.

Intuitively, Sam is not a cause of Victim's death in Disabling Samaritan. For, again, it seems that the death happened *despite* what Sam did, not *because* of what Sam did. However, in most cases of this sort it is likely that there will be features of the death that Sam caused. Imagine, for instance, that the way Assassin2's bullet pierced Victim's heart is slightly different from the way Assassin1's bullet would have pierced it. Or imagine that the amount of blood spilled is slightly different. Then Sam caused the death to involve a particular piercing of the heart, or a particular amount of blood spilled. Thus, it seems that, in Disabling Samaritan, Sam causes certain features of the death without causing the death itself.¹¹

Now, presumably, Disabling Samaritan is a case of causing only by preventing: presumably, for every feature of the death that Sam caused, she caused that feature only by preventing some other feature. First of all, there is a dependence relation between the fact that Assassin1 didn't shoot and the fact that Assassin2 did: in the circumstances, Assassin2's

¹¹ I am focusing on features concerning the manner of occurrence of the death, but, clearly, Sam can also cause features concerning its spatiotemporal location. Now, a question arises with respect to time. As I have pointed out, unlike delayers, hasteners typically strike us as causes. So, what if Sam actually *hastened* the death? (Suppose that Assassin2's gun works faster than Assassin1's, so that Victim ends up dying earlier rather than later given Sam's intervention.) In that case our intuitions might change. But, as I will suggest later, there is reason to be suspicious of our intuitions about hasteners.

shooting depended on Assassin1's not shooting. This dependence is clearly causal: in the circumstances, Assassin1's not shooting caused Assassin2 to shoot (the circumstances include the fact that Assassin2 had the disposition to shoot in case Assassin1 didn't and the fact that there were no obstacles to his shooting). Now, Sam wasn't causally responsible for the circumstances. Hence, Sam caused Assassin2 to shoot only by preventing Assassin1 from shooting. But all the features of the death that Sam caused were due to the fact that Assassin2 shot instead of Assassin1. Presumably, then, Sam caused the death to have those features only by preventing it from having other features (the features that the death would have had if Assassin1 had shot instead of Assassin2). As a result, the Causing by Preventing principle entails that Sam didn't cause the death.^{12, 13}

By the way, we can see under what conditions the principle *wouldn't* entail that Sam isn't a cause. Imagine that Sam was responsible for one of the circumstances at play. Imagine, for instance, that Sam carried a gun in her pocket, which she accidentally dropped while fighting with Assassin1, and that Assassin2 then used that gun to shoot Victim. In that case it wouldn't be true that Sam only caused by preventing: in that case Sam would also be causally responsible for one of the circumstances under which Assassin1's not shooting led to Assassin2's shooting. I think this is the right result: I think that in that case we would probably want to say that Sam *is* a cause of the death (although, of course, she is not morally responsible).¹⁴

¹² Again, the Kimean reasoning takes us as far as to the conclusion that Sam caused Assassin2 to shoot only by preventing Assassin1 from shooting. But it is natural to conclude from there that she also caused *the death* to have certain features only by preventing other features. For, if her only contribution to Assassin2's shooting was her causing Assassin1 not to shoot, and if Assassin2's shooting instead of Assassin1 accounted for the death's having certain features rather than others, then it seems that Sam's only contribution to the death's having certain features was her causing it to lack other features. (See also n. 9)

¹³ Recall that only ordinary properties of events are features (see section I). This restriction is important. Suppose, for instance, that we allowed "negative" properties of events as features. This would make trouble for the claim that Sam only caused features of the death by preventing other features. For, e.g., she caused the death *not* to involve the kind of piercing of the heart that only Assassin1's bullet produces, but, presumably, she didn't *do that* by preventing other features of the death.

¹⁴ Also, consider an "atypical" delayer as the one described in note 6: Assassin1 shoots a bullet at Victim; a second later, Assassin2 shoots a second bullet that deflects Assassin1's bullet and then knocks Victim down. This is an atypical delayer because it seems that Assassin2's shooting causes Victim's death, even if it delays it. I submit that Assassin2 doesn't cause only by preventing in this case. For Assassin2 doesn't

Cases like Disabling Samaritan are interesting because they have given rise to a controversy. Our intuition is that Sam doesn't cause Victim's death in that case. But this makes trouble for the (perhaps otherwise attractive) idea that causation is a transitive relation, i.e. for the idea that, if X causes Y, and Y causes Z, then X causes Z. For, even if Sam doesn't cause Victim's death, she clearly causes something that causes it: she causes Assassin2's shooting, which causes Victim's death. So transitivity seems to be violated in Disabling Samaritan.¹⁵

In light of cases of this sort, some philosophers believe that causation isn't transitive, but other philosophers prefer to bite the bullet, stick with transitivity, and reject our intuitions about those cases.¹⁶ Now, we have seen that there is a structure that apparent failures of transitivity have in common with other, less controversial cases (delayers): I have argued that they both exhibit the causing-only-by-preventing structure. I think that this provides at least a first step towards an argument against the transitivity of causation. For it is a reason to believe that apparent failures of transitivity are not an isolated phenomenon but an instance of a more general and pervasive phenomenon: the phenomenon of causing only by preventing. This is important because apparent counterexamples to transitivity have often been disregarded as exceptional or isolated cases by the friends of transitivity. In general, it might be methodologically warranted to disregard intuitions when they clash with more general principles we believe in, as long as they are, to some extent, isolated intuitions, and thus easily disposable ones. But now, in light of the concept of causing only by preventing, it starts to look as if our intuitions about apparent failures of transitivity are not as isolated as we thought they were.

contribute to the features of the death *just* by deflecting Assassin1's bullet: importantly, his contribution includes the shooting of the fatal bullet. In other words, Assassin2 is causally responsible for the circumstances under which the deflection of the first bullet leads to the impact between the second bullet and Victim.

¹⁵ Someone might try to argue that our intuitions about Disabling Samaritan are morally tainted, that is, that we are tempted to say that she is not causally responsible only because she doesn't bear any moral responsibility. However, this objection fails. Imagine a variation of the case where the interfering factor is not a moral agent. Imagine, for instance, that a powerful gust of wind blows Assassin1's gun away, and then Assassin2 steps in and kills Victim. This doesn't change our causal intuitions: it still seems that Victim's death occurs despite the wind blowing, not because of it.

¹⁶ For a defense of transitivity, see, e.g., Lewis (1986a), Hall (2000), and Paul (2000). For an attack on transitivity, see, e.g., Hitchcock (2001), Yablo (2004), and my (forthcoming).

In this section I have argued that Mackie's suggestion for delayers can be generalized to a principle that encompasses all features of events: the Causing by Preventing principle. I have also argued that the fate of this principle is tied to the fate of the claim that causation is a transitive relation. In the next section I examine in more detail the reasons for believing in the Causing by Preventing principle.

IV.

First, let me dispel a potential worry about the principle. As it turns out, Mackie is skeptical about the metaphysical prospects of her proposal for delayers. She claims that the proposal accounts for the way we think or talk about delayers, but that it might not reveal any deep metaphysical truth. More generally, Mackie is skeptical that we may use the notion of causing only by preventing to ground a genuine causal asymmetry between hasteners and delayers. This might be problematic because, although I haven't been concerned with arguing for an asymmetry between hasteners and delayers, any reason to doubt the general proposal about hasteners and delayers could be a reason to doubt the specific proposal for delayers, which I have relied on.

I will argue that Mackie is right to be skeptical about the general proposal but that we shouldn't thereby doubt the proposal for delayers. For the reason (at least, the best reason) to doubt the general proposal is *not* a reason to doubt the specific proposal for delayers.

Let me explain. As I have only briefly pointed out, Mackie provides an account for hasteners that complements her account for delayers. The account is: hasteners typically strike us as causes because they prevent only by causing, instead of causing only by preventing. This is to say, typically, a hastener prevents an event from happening later only by causing it to happen earlier (instead of causing it to happen later only by preventing it from happening earlier, which is what a delayer normally does).

Now, if one were to generalize this idea to other features of events, one would end up with one of the following principles:

Preventing by Causing (Strong): If, for *some* feature F that X causes an event E not to have, there is another feature G such that X causes E not to have F only by causing it to have G, then X is a cause of E.

Preventing by Causing Principle (Weak): If, for *every* feature F that X causes an event E not to have, there is another feature G such that X causes E not to have F only by causing it to have G, then X is a cause of E.

In other words, according to the Preventing by Causing principle, preventing by causing *is* causing (on the strong version, preventing *some* feature by causing some other feature is causing and, on the weak version, preventing *every* feature by causing some other feature is causing).

Now, there is good reason to believe that the Preventing by Causing principle is false on both of its versions. Arguably, it is possible to cause an outcome not to have a feature by making it have another feature *without thereby causing the outcome*. For instance, I can cause an explosion not to be red just by causing it to be blue, without causing the explosion itself. This could be true of all the features that I cause the explosion not to have (e.g., besides its not being red, its not being yellow, its not being green, etc.): I can cause the explosion not to have all those features just by causing it to be blue, without thereby causing the explosion. Similarly, I can cause a meeting not to be well attended just by causing it to be boring, without causing the meeting itself. Again, this could be true of all the features that I cause the meeting not to have (e.g., besides its not being well attended, its not being attended by fully awake people, etc.): I can cause the meeting not to have all those features just by causing it to be boring, without causing the meeting itself. This makes trouble for the Preventing by Causing principle on both of its versions.

Hence, whereas the Causing by Preventing principle is plausible, the Preventing by Causing principle is implausible. It is no mystery why this is so: it is harder to be a cause than to

not be a cause, and thus it is easier to give sufficient conditions for not causing than for causing. This suggests that there is reason to be skeptical about the prospects of using Mackie's proposal to ground an asymmetry between hasteners and delayers. For, if the generalized version of the principle for hasteners doesn't hold, then there is some reason to doubt the principle for hasteners itself. If preventing by causing is not always causing when the feature in question is something other than the time of occurrence, then why would it be causing when the feature is the time of occurrence? Why think that time is special in this way? Unless we see a reason to believe that time *is* special in this way, the fact that the principle doesn't generalize is a reason to be skeptical about its truth.¹⁷

I have argued that, although we probably should be skeptical about the hastener/delayer asymmetry (or about the prospects of grounding such an asymmetry on the causing-by-preventing/preventing-by-causing distinction), it doesn't follow that we should be skeptical about Mackie's principle for delayers, or about the Causing by Preventing principle in particular. For at least one main reason to be skeptical about the project as a whole isn't a reason to be skeptical about the subproject for delayers.

Now, all this shows is that a potential problem for the Causing by Preventing principle isn't really a problem. Of course, this doesn't by itself settle the issue of whether we should believe in the principle. So, what is the main reason to believe in the Causing by Preventing principle?

I think the main reason is something that I have already hinted at in the preceding sections. Namely, the principle captures the concept of something that, intuitively, contributes "only negatively" to an outcome. Call a thing of that type an "inhibitor." Intuitively, outcomes happen *despite* inhibitors, not *because* of them. The Causing by Preventing principle provides a way of cashing out the idea that inhibitors aren't causes. According to the principle, an inhibitor

¹⁷ I am not saying that time is *not* special in this way, but only that one would have to argue that it is in order to embrace Mackie's principle for hasteners; otherwise it would seem unmotivated.

causes only by preventing: to the extent that it is a cause, it is a cause only by virtue of being a preventer. This is so because an inhibitor contributes nothing to the outcome's having its actual features "over and above" what it contributes to its not having certain features. In other words, an inhibitor's contribution to an outcome is negative in that its contribution to the outcome's *features* is primarily negative in this way.

Recall that I take features to be ordinary properties of events, i.e. features are the sorts of things that determine what events are like. Thus, to say that an inhibitor contributes nothing to the outcome's having its actual features "over and above" what it contributes to its not having certain features is to say, basically, that it primarily determines what an outcome is *not* like, instead of what it *is* like. The suggestion is that it is in virtue of this primarily negative contribution that an inhibitor is not a cause of the outcome.¹⁸

Note that this account works even in a fully deterministic scenario. In an indeterministic scenario, it is tempting to suggest that inhibitors are things that lower the probability of an outcome, and to suggest that this is generally the explanation of why they aren't causes. As it has been pointed out, it is hard to make this idea plausible in the end.¹⁹ But a further problem with the probabilistic account is that the concept of an inhibitor still seems to make sense in a deterministic world. By accounting for the inhibiting relation in terms of the prevention of features instead of the lowering of probabilities, the Causing by Preventing principle allows us to make sense of inhibitors even in a purely deterministic scenario.

¹⁸ Since the Causing by Preventing principle is concerned with purely inhibitory relations, it fails to address cases of causing to happen in a certain way without causing to happen in which the contribution to the outcome is not negative but merely neutral. In section I, I mentioned two cases of this sort: causing an explosion to be blue without causing the explosion, and causing a meeting to be boring without causing the meeting. The Causing by Preventing principle only provides *sufficient* conditions for the absence of causation, and is silent about cases of this type.

¹⁹ For there seem to be counterexamples to the idea that, if I lower the probability of an outcome, I don't cause the outcome. Imagine that I start an unreliable process that cuts off an ongoing reliable process, but then the unreliable process goes to completion and issues in the outcome. Then I lower the probability of the outcome but I still cause it. For a more sophisticated attempt to capture the concept of an inhibitor in a probabilistic framework, see Eells (1991).

References

- Eells, Ellery (1991) *Probabilistic Causality*, Cambridge: Cambridge U. P.
- Hall, Ned (2000) "Causation and the Price of Transitivity," *Journal of Philosophy* 97, 4, pp. 198-222.
- Hitchcock, Christopher (2001) "The Intransitivity of Causation Revealed in Equations and Graphs," *Journal of Philosophy* 98, 6, pp. 273-299.
- Kim, Jaegwon (1974) "Noncausal Connections," *Noûs* 8, pp. 41-52, reprinted in *Supervenience and Mind*, Cambridge U. P., 1993.
- Lewis, David (1986a) "Causation," in *Philosophical Papers II*, Oxford U. P., pp. 159-213.
- Lewis, David (1986b) "Events," in *Philosophical Papers II*, Oxford U. P., pp. 241-269.
- Mackie, Penelope (1992) "Causing, Delaying, and Hastening: Do Rains Cause Fires?," *Mind* 101, pp. 483-500.
- Paul, Laurie (2000) "Aspect Causation," *Journal of Philosophy*, 97, 4, pp. 235-256.
- Sartorio, Carolina (forthcoming) "Causes as Difference-Makers," *Philosophical Studies*.
- Yablo, Stephen (2004) "Advertisement for a Sketch of an Outline of a Proto-Theory of Causation," in Collins, Hall, and Paul (eds.), *Causation and Counterfactuals*, MIT Press, pp. 119-137.