

Against Lewis's New Theory of Causation: A Story with Three Morals

Michael Strevens

Pacific Philosophical Quarterly, 84:4, pp. 398–412 (2003)

Abstract

A recent paper by David Lewis, “Causation as Influence”, provides a new theory of causation. This paper presents an argument against the theory, using a series of counterexamples that are, I think, of independent interest to philosophers of causation. I argue that (a) the relation asserted by a claim of the form “ C was a cause of E ” is distinct from the relation of causal influence, (b) the former relation depends very much, contra Lewis, on the individuation conditions for the event E , and (c) Lewis's account is unsatisfactory as an analysis of either kind of relation. The counterexamples presented here provide, I suggest, some insight into the reasons for the failure of counterfactual accounts of causal relations.

David Lewis's new account of causation (Lewis 2000) is, I believe, mistaken, but it is mistaken in instructive ways. In what follows, I present an argument against the account, based on a scenario that I claim constitutes a simple counterexample to all three of Lewis's counterfactual accounts of causation, and I draw several broad morals for the study of the truth conditions of claims of the form “ C is a cause of E ”.

1. Preliminaries: Concrete Events

Lewis’s account makes extensive use of the notion of what I call, after Hempel (1965, 421–3), a *concrete event*. A concrete event is an event that is individuated by every one of its intrinsic and spatiotemporal properties. To specify the concrete event of a particular jar’s shattering, for example, one would have to specify every detail of the shattering as finely as possible: the exact time that the shattering occurred, the exact number of pieces into which the jar shattered, the exact velocity of those pieces after the shattering, and so on.

Concrete events are, in Lewis’s terminology, extremely *fragile*: the smallest change in a concrete event “breaks” it, making it into a different concrete event. Consider the concrete event of the jar’s shattering. Had the jar shattered a moment later, or had it broken into a slightly different number of shards of glass, this concrete event would not have occurred; another concrete event, individuated by the new time or number of shards, would have occurred instead.

It seems that the events we have in mind when we make a causal claim such as “Sylvie’s throwing the rock caused the shattering of the jar” are, in many cases, not as finely individuated as concrete events. We accept that the event of the jar’s shattering could have occurred a moment later, meaning that very event, not some other, similar event. The concrete event of the shattering, then, is not the same thing as the event mentioned in the causal claim; it is, in some sense, more specific. Say that the concrete event *realizes* the event in the claim. Even if the events mentioned in causal claims do not have well-defined individuation conditions—and Lewis believes they do not—they do have well-defined concrete events as their realizers. It is these and other concrete events that figure in Lewis’s account of the truth conditions for causal claims.¹

1. What I call concrete events Lewis calls *alterations*. This has what I regard as the unfortunate consequence that the concrete event that realizes some event *C* is called an alteration of *C*, even though it changes nothing about the way that *C* actually occurred—in fact, in a sense, it *is* the way that *C* occurred. The term *concrete event* does not have this disadvantage, and has a long history in the literature. Lewis may, perhaps, use the new term because he wishes to keep open the possibility

2. Lewis's Account: The Simple Version

Let me begin (as does Lewis) with a simplified version of Lewis's new account. The counterexample will first be presented as an objection to this simple version; then, in section 4, it will be deployed against the official account.

According to the simplified version of Lewis's theory, a claim of the form " C was a cause of E " is true just in case, had C not occurred, the actual concrete realizer of E would not have occurred. Slightly less formally, C is a cause of E just in case, had C not occurred, either E would not have occurred, or E would have occurred but would have been realized differently.

To use Lewis's example, suppose that Sylvie and Bruno both throw rocks at a jar. Sylvie's rock arrives at the jar first and breaks it, but Bruno's throw was also fast and accurate, and had it not been for Sylvie's rock, Bruno's would have broken the jar. Thus, if Sylvie had not thrown her rock, the jar would still have broken, but its breaking would have been realized differently, because it would have occurred a moment later. (Recall that concrete events are individuated by, among other things, their exact time of occurrence.) It follows that Sylvie's throw is a cause of the breaking. Bruno's throw, by contrast, is not, because had Bruno not thrown his rock, the breaking would nevertheless have occurred in exactly the way and at exactly the time that it actually occurred.

Lewis raises a problem with the account as presented so far. Bruno's rock exerts a slight gravitational influence on Sylvie's rock. Thus, had Bruno not thrown his rock, the trajectory of Sylvie's rock would have been very slightly different. This appears to make Bruno's throw a cause of the breaking after all. Lewis says otherwise: very small differences do not count. Because the influence that Bruno's rock has on the manner of the jar's breaking is so slight, then, Bruno's rock does not count as a real cause.

Lewis goes on to make some comments about what counts as "very small".

of alterations that are not as finely individuated as concrete events. This possibility will play no role in what follows.

These comments make no difference to what follows, so I put them aside and move on to the argument.

3. A Counterexample

Consider the case of Sylvie and Bruno with the following amendments. First, the implements of destruction are not rocks but something more elastic, that is, something more inclined to bounce, say, steel balls. Second, and more importantly, Sylvie's ball hits the jar and Bruno's doesn't, not because Sylvie's ball gets there first, but because the balls collide in mid-air, and as a result, Bruno's ball is deflected so as to miss the jar altogether.

The relation between the trajectories of the two balls is, I will suppose, as shown in figure 1. The solid line represents the trajectory of Sylvie's ball, the




Figure 1: The trajectories of Sylvie's and Bruno's balls

dashed line the trajectory of Bruno's ball.

If Sylvie's and Bruno's balls hit one another squarely, they in effect exchange trajectories. The trajectory of Sylvie's ball after the collision is identical to the trajectory that Bruno's ball would have taken if the collision had not occurred, and vice-versa. What's more, if the balls are traveling at the same speed and have the same mass, Sylvie's ball after the collision will have the same speed as Bruno's ball before the collision. It follows that, if Sylvie had not thrown her ball, the jar would have been struck by Bruno's ball at exactly the same time and in exactly

the same place as it was in fact struck by Sylvie's ball.² Assuming that the balls are identical in other respects, the jar would have broken in exactly the same way as it actually did. That is, the counterfactual breaking would have exactly the same realizer as the actual breaking.

According to Lewis's account, then, Sylvie's throw is not a cause of the jar's breaking. But this is clearly incorrect. It is Sylvie's ball that smashed the jar, and for this reason, Sylvie's ball is a cause of the breaking. This is not to say that Bruno's throwing is not also a cause; in fact, it probably ought to be counted as such. The counterexample does not turn on the status of Bruno's throw but on the status of Sylvie's: Lewis's new account incorrectly denies that Sylvie's throw is a cause of the breaking.

The scenario is, of course, a special case of the kind of preemption scenario that created problems for Lewis's early account of causation, according to which C is a cause of E just in case, had C not occurred, E would not have occurred (Lewis 1973a). Here I have set things up so that, not only is there a backup cause that would have caused E even if C had not, but so that the backup cause would have caused E in exactly the same way that C actually did. I believe, for reasons given in section 6, that preemption cases reveal a fundamental and ineliminable flaw in counterfactual accounts of causation.

Lewis's early counterfactual account was rescued from some kinds of preemption objections by an appeal to chains of dependence (Lewis 1986c). The idea is that, although preemption may nullify the counterfactual dependence of an effect E on its legitimate cause C , some intermediate event D might exist, such that E depends counterfactually on D , and D depends counterfactually on C . The exist-

2. The physics here is somewhat idealized: most notably, in the main text and in the figure, I am treating the balls as point particles. Since the balls in fact have some diameter, an exchange of trajectories will leave Sylvie's ball striking the jar slightly to Sylvie's side of where Bruno's ball would have struck the jar. But this complication can be dealt with by putting a slight spin on the balls, so as to alter the direction of Sylvie's ball after the collision in such a way as to make up for the diameter effect.

tence of such a chain of dependence is then declared to be sufficient for C to count as a cause of E . Can Lewis's new account be saved from my counter-example in the same way?

The attempted solution would, of course, involve an amendment to the new account, so that a chain of dependence rather than dependence itself is sufficient for causation. Lewis's new account is, as he presents it, designed in part to avoid the need to appeal to chains in solving preemption problems, so to make the amendment is to admit at least partial defeat.³ But never mind; does the amendment work?

It seems to me that it does not. Consider first the obvious choice of an intermediate event D , the collision between Sylvie's and Bruno's balls. To dodge the counterexample, it must be shown that Sylvie's throw is a cause of the jar's breaking, and thus that (a) the collision depends on Sylvie's throw, and (b) the jar's breaking depends on the collision. On Lewis's new account, (a) is true but (b) is not. If Sylvie had not thrown, the concrete realizer of the collision would not have occurred (indeed, there would have been no collision), but had there been no collision, the concrete realizer of the jar's breaking would still have occurred.⁴

3. Note that Lewis nevertheless resorts to chains to solve certain problems later in the paper; see section 4.

4. Here I assume that the closest possible world in which the balls do not collide is one in which they have the same trajectories up until the moment of collision, at which time a Lewisian "small miracle" occurs, and the balls pass right through one another. Bruno's ball therefore flies on to break the jar. This is, I think, the evaluation of the counterfactual mandated by Lewis's account of closeness, on which the closest world is one that matches the facts of the real world—in particular, the positions of balls at times—as closely as possible up to the time of the occurrence of the antecedent of the counterfactual (Lewis 1973b, 1979).

A referee for this paper suggested that our intuitions about counterfactuals may be fuzzier than Lewis's account allows. For the purposes of evaluating the counterfactual, it might be reasonable to imagine that Bruno's ball swerves at the last moment, rather than passing right through Sylvie's ball. If so, then it is not true that, had the collision not occurred, the jar would still have broken in exactly the way it did. It is true, however, that had the collision not occurred, the jar *might* still have

Now consider an event D that occurs after the collision, namely, the “event” of Sylvie’s ball being at such and such a place and traveling at such and such a velocity just prior to the shattering. Does the shattering depend on this event? On Lewis’s scheme for evaluating counterfactuals it does. But note also that, in Lewis’s metaphysics, the concrete realizer here is not *Sylvie’s* ball’s having such and such a velocity and position at such and such a time, but a ball with such and such intrinsic properties (weight, composition, and so on) having said position and velocity at said time. Assume that Sylvie’s and Bruno’s balls are identical. Then the event of such a ball having said position and velocity at said time does not depend on the collision, since, had the collision not occurred, there would have been a ball with all the stated physical properties at the stated time, namely, Bruno’s ball.

The collision, then, is the weak link in the chain of counterfactual dependence: the shattering of the jar depends on events (namely, ball states) after the collision, and the collision depends on events before the collision, but events after the collision do not depend on the collision, or any event before the collision, since had there been no collision, a physically identical chain of post-collision events would anyway have led up to the shattering of the jar. The ball-throwing scenario, then, provides a single counterexample to Lewis’s early, middle, and new accounts of causation. Or it least, it does if it carries over from the simple to the official new account. To the official account I now turn.

broken in the way it did, since the passing-through envisaged in the last paragraph is one way that the collision might not have occurred. It follows, on the standard semantics for counterfactuals, that it is false that, had the collision not occurred, the jar would *not* have broken in the way it did. Consequently, the necessary dependence relation between the collision and the breaking does not hold, and the counterexample stands.

4. Causation as Causal Influence

Lewis's official new theory of causation is built around the notion of causal influence. According to Lewis, the degree of causal influence that one event C has on another event E can be assessed as follows. Take the concrete event that realizes C ; call it c . Make a small change in c so as to produce a different concrete event c' . Assess the corresponding change in the concrete realizer of E . That is, if e is the actual concrete realizer of E , ask: had c' occurred instead of c , how would e have been different? More precisely, what concrete event e' , if any, would have occurred instead of e ? Repeat for many different variations of c . Then the degree of causal influence that C has on E is roughly proportional to the average degree of difference between e' and e .⁵ If varying c makes no difference to e —if e occurs no matter how c is varied—then C does not influence E at all. If varying c makes a big difference—if varying c usually results in the occurrence of concrete events quite unlike e —then C influences E strongly. And so on.

What is the relation between causal influence and causation? On p. 191 of his paper, Lewis writes that “event C causes event E if and only if there is a chain of dependencies running from C to E ”, where it is clear from the context that the dependencies in question are relations of causal influence. This is not quite consistent with the comments made elsewhere about the proper handling of, for example, the gravitational influence of Bruno's ball on Sylvie's ball, which discount small influences as causes. I will interpret Lewis's formulation as follows. C is a cause of E just in case either

1. C has a relatively large influence on E (relative, that is, to the influence of the other causal influences on E), or
2. C has a relatively large influence on some intermediate event D which is itself a cause of E .

5. Lewis has a little more to say about the measurement of influence, but his comments will not be important here. See, however, section 5.3 for an alternative method of measuring influence.

This recursive definition, then, identifies C as a cause of E just in case there is a chain of events beginning with C and ending with E , such that each event in the chain has a sufficiently large influence on the next. In other words, causation is roughly the ancestral of causal influence—“roughly” because of the requirement of relative magnitude. Call this the *causation equals influence* account of the truth conditions for causal claims. Lewis’s official new theory of causation has two parts: the *causation equals influence* account of causal claims, and an account of causal influence.

I will take this opportunity to argue against both parts of Lewis’s theory, first, the equation of causation and (chains of) causal influence, and second, the account of causal influence itself. The scenario presented in section 3 can be used to make both arguments. In a certain sense, however, it would be misleading to use it, or any other example of preemption, to contest the *causation equals influence* claim, because what is wrong with this claim at root has nothing to do with preemption. For this reason, my argument against *causation equals influence* in this section will use a different example that is not a case of preemption.⁶ In section 5, I will return to the preemption case to argue that Lewis’s account is not a good theory of causal influence.

Suppose that Sylvie fires a pellet from her laser-sighted slingshot at a dartboard. Although she does not know it, Bruno has sabotaged her laser sight so that the slingshot consistently shoots too low. As Sylvie shoots, a strong lateral gust of wind springs up from her left. The wind blows the pellet to the right; meanwhile, because of Bruno’s act of sabotage, the pellet is aimed too low. The combined result of these two perturbations is that the pellet hits a spot below and to the right of the bull’s-eye. Now: what is the cause, or what are the causes, of the pellet’s hitting the dartboard too low? The correct answer: Bruno’s sabotage is a cause, but the gust of wind, which was responsible only for the pellet’s sideways drift, is

6. For an argument against the *causation equals influence* claim that does use a case of preemption, see Schaffer (2001).

not a cause.⁷

Lewis's account fails to reproduce this answer correctly, however. It holds that the gust is a cause of the pellet's landing too low. To see this, consider the concrete realizer of the event of the pellet's landing too low. This is the concrete event of the pellet's hitting such-and-such a point on the dartboard at such-and-such a time (but the time will be unimportant in what follows). Varying the wind makes just as much of a difference to this concrete event as varying the angle of Sylvie's laser sight. (I have, of course, gone out of my way to construct an example where the difference made by each factor is of exactly the same sort, so that questions of how to weight different kinds of variation in a concrete event do not arise.) Thus

1. The wind is just as great a causal influence on the shot as the sabotage, so
2. It follows from the *causation equals influence* doctrine that the wind has the same claim as the sabotage to be a cause of any event for which the shot is a concrete realizer.

Hence the wind causes the shot to go too low.

The erroneous step appears to be (2) rather than (1). Intuitively, the wind really does causally influence the shot. But it does not cause the shot to go too low. Where Lewis goes wrong is in equating the causal relation between events and the causal influence relation between their concrete realizers.

More exactly, on Lewis's account, an event C 's degree of causal influence on another event E depends only on E 's concrete realizer, and not on E itself. Thus,

7. Is the fact of the pellet's hitting the dartboard too low an event? Or, what is more important, is it a legitimate relatum of the *is a cause of* relation? By Lewis's lights, yes: any state of affairs can be a cause or can be caused (Lewis 1986a). Furthermore, it is quite normal to talk of the causes of such things as the pellet's landing too low (think of a conversation during artillery practice).

If necessary, however, the example could be changed, at a small cost in complexity, so that the effect in question is a more traditional causal relatum. Suppose that below the bull's-eye lies a wide window; then it seems right to say that the sabotage, but not the gust, is a cause of the pellet's hitting the window. The objection stated below will still go through.

on the *causation equals influence* account, if an event is a cause of one event with a certain concrete realizer, it is a cause of all events with that realizer. But this is not the way we think of causation. The sentence “*C* is a cause of *E*” does not necessarily remain true when *E* is replaced by another event having the same realizer. The reason, it seems, is that the high level event *E* picks out a certain property *P* of the realizer—the property in virtue of which the realizer counts as a realization of *E*—and the causal claim asserts not just that *C* has an effect on some property or other of the realizer, but that it has an effect on *P*.

If my diagnosis is correct, then Lewis has made a wrong step very early on in his new approach to causation. Citing his belief that the high level events in causal claims are rather indeterminate entities, and that the identity of their realizers may be one of the few determinate facts about them, he has resolved to give an account of causal claims that adverts only to facts about realizers. This now appears to be a mistake: it implies that causal claims remain true under a certain substitution operation, that of replacing a high level event with a different high level event having the same actual concrete realizer, when in fact they do not.

5. Causal Influence Alone

5.1 *The Counterexample Again*

Never mind causal claims; is Lewis’s account of causal influence a good one? Before giving an answer to this question, two caveats. First, Lewis may have intended his account of causal influence as a piece of technical machinery, not as an account of an everyday notion. Second, it is not obvious that we use the term *causal influence* in a univocal way.

In answer to the second worry, observe that there are some claims about causal influence that are uncontroversially true or uncontroversially false, regardless of the context in which they are made. For example, it is false that your astrological birth sign causally influences your life (unless, of course, you act on astrologically motivated advice), whereas it is true that the moon causally influences the sea level

in any given locale. It is worth asking, I think, whether Lewis's account of causal influence does a reasonable job of capturing the truth conditions for claims such as these, regardless of whether or not this was his intention. Needless to say, my focus will be on examples of less astronomical proportions.

Consider the scenario proposed in section 3, where Sylvie's and Bruno's balls collide on the way to the jar. At first, Lewis's account of causal influence looks promising. Varying the details of Sylvie's throw will indeed make a difference to the exact manner of the breaking of the jar. Thus Sylvie's throw, whether or not it counts as a cause of the breaking, will count as a causal influence on the breaking.

But oddly, this is only true for small variations in Sylvie's throw: large variations will result in Sylvie's and Bruno's throws no longer colliding, in which case Bruno's throw will break the jar in exactly the way that Sylvie's throw actually broke it, so there will be no difference between the counterfactual and the actual concrete realizers of the breaking. Thus small variations in Sylvie's throw will make a difference to how the breaking is realized, but larger variations will not.

This peculiar behavior opens the door to a refinement of the counterexample to deal with the move from cause to causal influence. Suppose that Sylvie's and Bruno's balls are very small and moving very fast. Perhaps they are fired from guns. Then even very small variations in Sylvie's "throw" will result in no collision, hence no change in the realization of the breaking. By changing the example so that the balls are smaller and faster, one can reduce the average variation corresponding to, thus the degree of causal influence of, Sylvie's throw to as near zero as you like. This seems wrong.⁸

Worse, you cannot perform the same trick with Bruno's throw. Even if there is no collision, variations in Bruno's throw will bring about variations in the break-

8. Note also that this constitutes a counterexample to Lewis's account of causal claims. Sylvie's throw counts as a cause of the jar's shattering, recall, just in case it has a sufficiently large causal influence on the shattering. But using the suggested technique, the influence of the shattering can be reduced to the point that the throw no longer counts as a cause—even though it is Sylvie's ball that shatters the jar.

ing. Thus, at the very least, Bruno's throw will turn out to have a much greater degree of causal influence on the breaking of the jar than Sylvie's throw. But in the actual case—the case where it is Sylvie's ball that strikes the jar—this seems quite incorrect. Lewis's account not only gets the truth conditions for our causal claims wrong, but also gives the wrong truth conditions for our claims about degrees of causal influence.

Let me consider three responses on Lewis's behalf. The first two responses attempt to show that Lewis's account of causal influence does, after all, count Sylvie's throw as being just as influential as Bruno's. The third accepts that Lewis's account classifies Sylvie's throw as less influential, but argues that this classification is correct.

5.2 *The Appeal to Transitivity*

Can Lewis appeal to transitivity? In section 3, I considered the possibility that Sylvie's throw is a cause of the jar's shattering because it is a cause of an intermediate event which is in turn a cause of the shattering. Lewis's official account defines causation as the ancestral of influence. Thus on this account, Sylvie's throw would count as a cause of the shattering even if it did not causally influence the shattering, provided that it causally influenced some intermediate event that in turn causally influenced the shattering. Can a similar maneuver save Lewis's account of causal influence? Apparently not. It would be peculiar to define *causal influence* as the ancestral of causal influence, and in any case, Lewis holds that the relation of causal influence is not transitive (Lewis 2000, 191).

Still, it is worth demonstrating that transitivity will not help here, for much the same reasons that it did not help the simple account in section 3. Suppose we interpret Lewis's definition of causal influence as a definition of what I will call *direct* causal influence, and then we define causal influence as the ancestral of direct causal influence.⁹

9. Of course, given the continuity of physical processes, *direct* here cannot mean *unmediated*.

In order to save Lewis, we need to find an event D between Sylvie’s throw and the jar’s shattering such that (a) Sylvie’s throw causally influences D to a reasonably large degree (that cannot be reduced by making the throw faster or Sylvie’s ball smaller), and (b) D causally influences the shattering to a reasonably large degree.

Is the collision itself such an event? No, because making the ball smaller and faster reduces the collision’s influence on the jar’s shattering to as close to zero as you like, since, when the colliding balls are sufficiently fast and small, very small changes in the collision’s concrete realizer will result in the balls’ not colliding at all,¹⁰ in which case Bruno’s ball will take the place of Sylvie’s, and the chain of events leading up to the jar’s shattering will be physically identical to the actual chain of events.

For essentially the same reason, no other event can serve as the required intermediary: the influence of the throw on events after the collision can be made arbitrarily small, as can the influence of events before the collision on the breaking. Thus there is no event that both influences the breaking to a sufficient degree, and is itself influenced by the throw to a sufficient degree.

5.3 *Measuring Causal Influence*

Jonathan Schaffer has suggested to me the possibility of an alternative interpretation of Lewis’s method for measuring causal influence between two events C and E . Rather than altering the concrete realizer c of C by a fixed amount, and tallying the *average* change in the concrete realizer e of E , Schaffer suggests tallying the

Consider it a term of art.

10. There is a temptation here to confuse “small changes in the realizer” with “small changes in the collision”, the latter expression implying that the changes are restricted so as not to alter the fact that there is a collision. But “small changes in the collision” play no part in Lewis’s theory. Lewis’s theory is designed to attend only to realizers, and not to depend at all on the individuation conditions for a high level event such as the event of the balls colliding. In particular, what it is to count as a “small change” to a realizer ought not, if we are to respect Lewis’s intentions, to depend on the individuation conditions for any high level event.

maximum change in e . On this measure, it is irrelevant to the degree of influence that most changes to c have no effect on e . All that matters is the greatest change in e that can be induced by making some change to c . In particular, then, if altering c almost always has no effect on e , but there is one, very improbable, change in c that does have a great effect on e , then C counts as a strong influence on E .

This measurement scheme would seem to have the desired effect: Sylvie's throw will count as a strong causal influence on the jar's shattering, because although most changes to the throw have no effect on the shattering whatsoever (because there is no collision), there exist a small set that alter the collision in such a way as to shatter the jar in a number of different ways (because the collision occurs, but somewhat differently from the way it occurs in the actual case).

This solution to the problem works because it ignores the relative probabilities with which variations bring about effects. In so doing, it opens Lewis's account up to a different kind of counterexample. Suppose that, at the same time that Sylvie and Bruno are out throwing rocks, their mother is hunting sheep with a high-powered rifle. Just after Sylvie and Bruno throw their rocks, Mother fires. She is miles away from the jar, and aiming in a different direction; consequently, her bullet goes nowhere near the jar. But suppose that, if Mother's aim had been slightly off, the bullet could have hit a series of rocks, ricocheting from one to the next, in such a way that it struck the jar. There is an alteration of the concrete realizer of Mother's shot, then—an absurdly improbable alteration—on which the shot shatters the jar. Furthermore, realizers in a very small neighborhood of this one also presumably shatter the jar, but in rather different ways. On the new measuring scheme, then, Mother's shot counts as a considerable causal influence on the shattering of the jar, even though the bullet goes nowhere near the jar and (perhaps more importantly) was overwhelmingly likely not to.

This example, though contrived, is of a kind that is presumably common in everyday life: although we never do know it, it is often the case that some minor event C might have had a noticeable consequence E had C been realized slightly differently. Yet in those cases where E occurs for some other reason, we certainly

do not count C as a causal influence on E .

It seems, then, that the attempt to make Lewis's account of causal influence sensitive enough to attribute the proper causal import to Sylvie's throw makes it far too sensitive: it attributes import to events that, intuitively, have little or no influence at all. The problem, I think—here and elsewhere—is in the appeal to counterfactual possibilities, for more on which see section 6.

5.4 *Is Sylvie's Throw Really as Influential as Bruno's?*

Consider the following last line of defense: Sylvie's throw really ought not, after all, to be counted as so great a causal influence on the breaking as Bruno's. The reason for holding this view is the following asymmetry: had Sylvie not thrown her ball, the jar would have been broken anyway, but had Bruno not thrown his, the jar would not have been broken. Lewis's account of causal influence, it might be claimed, simply captures this causal asymmetry.

Perhaps the most satisfying response to this defense is to alter the counterexample so that the asymmetry is removed, that is, so that Bruno's throw is no longer required for Sylvie's throw to break the bottle. So far I have assumed that Sylvie's throw was rather badly aimed. But suppose that Sylvie was in fact making a tricky bank shot: had her ball not collided with Bruno's, it would have ricocheted off several hard surfaces and broken the jar after all (see figure 2). The breaking would




Figure 2: Sylvie's attempted bank shot

have been realized differently from the actual breaking, because—given the complicated trajectory of the bank shot—it would have occurred later, and the ball would have been coming from another direction. But it would have occurred all the same. Note that, since the balls exchange trajectories, Bruno’s ball follows the intended trajectory of Sylvie’s ball. But by the time it gets to the jar, the jar has already been shattered by Sylvie’s ball.¹¹

In this new scenario, Sylvie’s throw would have broken the jar whether Bruno had thrown or not. Thus the asymmetry that provided a reason for thinking that Sylvie’s throw was a less important influence than Bruno’s is removed. The remarks made about degrees of causal influence within the original scenario also hold within the new scenario. Small variations in Sylvie’s throw make small variations in the breaking. But larger variations—large enough to avoid the collision—make no difference to the way the breaking occurs, since Bruno’s ball will reach and break the jar before Sylvie’s ball, even if it makes the bank shot, gets there. Thus Lewis’s account will, as in the original scenario, count Bruno’s throw as a much greater influence on the breaking than Sylvie’s, asserting the existence of a causal asymmetry that is simply not there.

6. Conclusions

I draw the following morals for the study of the truth conditions for causal claims. First, as shown by the example of Bruno’s sabotage of Sylvie’s crossbow in section 4, causation and causal influence are distinct relations.

Second, Lewis’s attempt to provide truth conditions for causal claims that depend in no way on the individuation conditions for the high level events or states of affairs connected by the claims is misguided. In the example given in section 4, it is impossible to determine the cause of Sylvie’s pellet’s landing too low without taking into account what aspects of that event’s concrete realizer make it an in-

11. This scenario provides a case of what Lewis calls *late cutting* or *late preemption*.

stance of the pellet's landing too low, by contrast with, for example, those aspects that make it an instance of the pellet's landing too far to the right. Intuitively, an influence on an event E 's concrete realizer e can only be a candidate cause for E if it influences those aspects of e in virtue of which it counts as a realizer of E .

This intuition suggests that being a cause of E has something to do with making a difference to E 's occurrence. But third, I submit that the counterexamples deployed in this paper provide further evidence that the appropriate difference-making relation is *not* the relation captured by counterfactual claims of the form: "had C not occurred, E would not have occurred". Rather, it is some other difference-making relation; call it, for the sake of the next few paragraphs, *causal difference-making*.

How does counterfactual difference-making differ from causal difference-making? The identity of the causal difference-makers depends, I propose, only on the actual causal processes that lead to E . (I am helping myself here to a notion of a causal process that is prior to the notion of causal difference-making.) The identity of the counterfactual difference-makers depends, by contrast, also on causal processes (often non-actual) that might have led to E , but did not. The elements of these processes, then, might have been causal difference-makers, but were not. For this reason I call them *would-be difference-makers*. The problem with the counterfactual criterion for causation is that it takes these would-be difference-makers into account, and in doing so, allows them to mask the actual difference-making power of the real causes. Hence the counterfactual account's preemption problem.

The history of the counterfactual account of causation consists of attempts to render irrelevant the would-be difference-makers. This is normally done by specifying some details of the chain of events leading to E (or in Lewis's new account, by stating some details about the realization of E that could only have been brought about by a very specific chain of events), and by holding that, in order to be relevant, a factor must not merely stand to make a difference to E , but must stand to make that difference in precisely the specified way. In most cases, the would-be difference-makers would have made their difference in some other way; they are

therefore ruled irrelevant, leaving the account to pick out the real causes of E as though the would-be causes had not been present.

The history of attempts to refute the counterfactual account may then be seen as an effort to construct counterexamples in which the would-be difference-makers come as close as possible to the actual difference-makers in their mode of difference-making. Cases of late preemption (see section 5.4), trumping (Schaffer 2000), and, of course, the counterexamples offered in this paper, can all be regarded in this way.

If my comments on the link between difference-making and causation are on the right track, then this arms race can never reach a satisfactory conclusion. The reason that would-be difference-makers are irrelevant to determining the causes of an event is not, as amendments to the counterfactual account implicitly suppose, that would-be difference-makers make their difference, *subjunctively*, in the wrong sort of way. It is rather that they *actually* make no causal difference at all.

The correct criterion for causal difference-making will not, then, be a counterfactual criterion. A counterfactual criterion will never show the disdain for unactualized possibilities that is required of the true criterion. But the true criterion will be a criterion for difference-making, all the same. One conception of difference-making that might be suitable is described in Strevens (2003). This paper will, I hope, inspire the search for others.

References

- Hempel, C. (1965). *Aspects of Scientific Explanation*. Free Press, New York.
- Lewis, D. (1973a). Causation. *Journal of Philosophy* 70:556–67.
- . (1973b). *Counterfactuals*. Harvard University Press, Cambridge, MA.
- . (1979). Counterfactual dependence and time’s arrow. *Noûs* 13:455–76.
- . (1986a). Events. In Lewis (1986b).
- . (1986b). *Philosophical Papers*, volume 2. Oxford University Press, Oxford.
- . (1986c). Postscript to “Causation”. In Lewis (1986b).
- . (2000). Causation as influence. *Journal of Philosophy* 97:182–97.
- Schaffer, J. (2000). Trumping preemption. *Journal of Philosophy* 97:165–181.
- . (2001). Causation, influence, and effluence. *Analysis* 61:11–19.
- Strevens, M. (2003). The causal and unification accounts of explanation unified—causally. Forthcoming, *Noûs*.