

17 Do All and Only Causes Raise the Probabilities of Effects?

Christopher Hitchcock

1 Introduction

According to probabilistic theories of causation, causes raise the probabilities of their effects. Opponents of probabilistic theories of causation offer putative counterexamples. A moment's reflection should lead us to expect such counterexamples to be of two types: (1) causes that appear *not* to raise the probabilities of their effects; and (2) events that appear to raise the probabilities of other events, without causing those events. Almost all of the cases that have been discussed in the literature have been of the first sort; these can be effectively handled using resources that are already available. Counterexamples of the second sort, which have been largely ignored, still pose a threat. I will explore some options for avoiding these counterexamples, and more important, argue that they raise fundamental questions about the nature of indeterministic causation—questions that transcend issues about the correctness of any particular philosophical theory.

2 Probabilistic Causation

The motivation for a probabilistic theory of causation lies in the desire to have a theory that encompasses causal relationships that are inherently indeterministic or “chancy.” The central idea behind probabilistic theories of causation is that causes raise the probabilities of their effects. Perhaps the simplest and most natural way to formulate this idea is in terms of conditional probability. C raises the probability of E , and hence C causes E , exactly if $P(E/C) > P(E/\sim C)$.

This proposal is subject to a well-known problem. Coughing raises the probability of lung cancer: Coughers are more likely to be smokers than are noncoughers, and since smokers are more likely to contract lung cancer than nonsmokers, coughers will also be more likely to contract lung cancer than noncoughers. Coughing does not, however, *cause* lung cancer—the correlation between coughing and lung cancer is *spurious*. Spurious correlations can give rise to counterexamples of both sorts described in the introductory section.

The standard solution to this problem is to require that causes raise the probabilities of their effects *ceteris paribus*. There are at least two well-developed approaches based on this idea. According to one (see, e.g., Cartwright 1979; Skyrms 1980; Humphreys 1989; Eells 1991), C causes E iff $P(E/C\&B) > P(E/\sim C\&B)$, for a variety of background conditions B . By conditioning on background conditions, we hold fixed the presence or absence of confounding factors.

Lewis (1986) analyzes the concept of *ceteris paribus* probability-raising differently. According to his theory, *C* causes *E* if the probability that *E* would occur in the actual world, at the time when *C* occurred, is higher than the probability that *E* would occur at the corresponding time in any of the closest possible worlds in which *C* does not occur.¹ In this theory single-case probabilities in different possible worlds, not conditional probabilities, are compared. For example, suppose that Barney smokes, coughs, and develops lung cancer. Barney's coughing is a cause of his lung cancer if he would have been less likely to contract lung cancer had he not coughed. But the relevant counterfactual is false: The closest possible world in which Barney does not cough is one in which he nonetheless smokes. In such a world, the probability of his developing lung cancer remains unchanged. In what follows, I will not distinguish between the different methods of articulating the basic probability-raising idea unless necessary.

My focus will be on what I call "actual causation": the relation that holds between two particular events when one does, in fact, cause the other. This is to be distinguished from the relation that holds when one event has (or is of a type that has) a *tendency* to cause another.² As we shall see, the most troubling sort of counterexample to probabilistic theories of causation threatens only theories of *actual* causation. Even if the counterexamples cannot be circumvented, probabilistic theories of causal tendencies remain attractive. Unlike some authors (e.g., Sober 1985; Eells 1991), I find it unpalatable that these are simply two independent causal relations: Surely whether one event did in fact cause another has a great deal to do with whether the one (was of a type that) had a tendency to cause the other (see Hitchcock 1995a for further discussion). This leaves us with good reason to continue pursuing probabilistic theories of causation even in the face of the difficulties discussed below.

3 Making It the Hard Way

Many authors have argued that there are examples in which a cause lowers the probability of its effect, even when one holds other factors fixed. Perhaps the most famous example, due originally to Deborah Rosen, is presented in Suppes (1970). (See also Salmon 1984, chapter 7.) A golfer lines up to drive her ball, but her swing is off and she badly slices the ball, sending it on a trajectory well to the right of the hole. Her slice decreases the probability that it will land in the cup for a hole-in-one. By chance, however, the ball bounces off a tree trunk at just the right angle to send it on a trajectory back toward the cup. As it happened, her slice did cause the ball to go into the cup, even though the slice lowered the probability of this outcome.

This sort of example has received considerable attention. I will present only the solution that I favor. Probabilistic theories of causation require that we compare the probability of the effect in the presence of the cause, and also in its absence. But a cause may be absent in many ways. For example, our golfing heroine might have hit the ball squarely, hooked the ball to her left, or barely nicked the ball (causing it to move hardly at all); her swing might have been so bad that she missed the ball completely, or she might have refrained from swinging. These are all ways in which she might not have sliced the ball. The intuition appealed to in the previous paragraph, that the slice lowered the probability that the ball would land in the cup, suggests that the relevant alternative is the one where she hits the ball squarely. But why should this be so? Perhaps the relevant alternative is the one in which the golfer refrains from swinging altogether; relative to such an alternative, the slice actually increases the probability of a hole-in-one. My view (developed in Hitchcock 1993, 1996a,b) is that there is no objectively correct alternative for purposes of probability comparison. Rather, causal claims are contrastive in nature; they are true or false relative to a specific alternative. Thus, the golfer's slice caused the hole-in-one, relative to the alternative in which she abstains from swinging, but not relative to the alternative in which she hits it squarely. In the latter case, we say that the ball landed in the cup *despite* the badly sliced shot.

There are a number of advantages to this approach. First, it avoids the need to select a privileged alternative from among the many alternatives to the candidate cause. Second, it can accommodate the seemingly incompatible judgments that "the golfer's slice caused the hole-in-one" and "the ball landed in the cup *despite* the badly sliced shot" are both true causal claims. We can accept both claims because we can hear them (perhaps with the help of a pragmatic rule of accommodation) as being made relative to different alternatives. Third, the relativity of causal claims to specific alternatives can account for a curious linguistic feature of causal claims. To borrow an example from Dretske (1977b), suppose that Susan has stolen a bicycle, resulting in her arrest. We might agree that Susan's *stealing* the bicycle caused her to be arrested, while denying that Susan's stealing *the bicycle* caused her to be arrested. The placement of stress in a causal claim can affect its truth value. According to the proposed view, the different placement of stress points to different alternatives. Susan's stealing the bicycle caused her to be arrested relative to alternatives in which she acquired a bicycle by some other means, but not relative to alternatives in which she stole some other item.

I conclude that putative counterexamples involving probability-lowering causes can be eliminated using a refinement to probabilistic theories of causation that is motivated on a number of independent grounds.

4 Three Scenarios

Peter Menzies (1989a, 1996) describes a case involving causal preemption in which a probability-raising event is not a cause. He takes such cases to be isolated and atypical:³ “It is striking how a probabilistic theory of causation, appropriately formulated, follows the contours of our intuitive conception of causation with amazing accuracy—with the exception of the problem cases to do with pre-emption and over-determination” (Menzies 1996, p. 100).

In fact, however, preemption is in no way essential to Menzies’s example. James Woodward (1990) describes a case with the following structure. Suppose that C_1 and C_2 both increase the probability of E , and would do so regardless of whether the other is present or not. On some particular occasion, C_1 , C_2 , and E occur. Although both C_1 and C_2 are potential causes of E , could it not be the case that in fact only C_2 causes E on this particular occasion? That is, couldn’t this be one of those cases where C_1 fails to cause E , not because it is preempted, but simply because its operation is chancy? If so, then C_1 raises the probability of E without causing it.⁴

Woodward offers the following more specific example in a review of Humphreys (1989):

Scenario 1: Suppose we know ... that each of two different carcinogenic materials C_1 and C_2 [e.g., gyromitrin and diazonium metabolite] can cause [stomach tumors] in mice (E). Suppose also that the operation of each of these causes is non-deterministic but governed by a stable probability distribution: each increases the probability of cancer, although to some value strictly less than one. Suppose also that there is no evidence for any interaction effect between C_1 and C_2 when both are present. Now suppose that a particular mouse is exposed to ... C_1 and C_2 and develops cancer E . It follows on Humphreys’ account that since both C_1 and C_2 increase the possibility of cancer, both cause or have causally contributed to cancer. But why should we believe this? How do we know that the cancer was not instead caused by C_1 alone or C_2 alone? We know that when ... C_1 occur[s] in isolation (without ... C_2) ... , it is perfectly possible for [it] to increase the probability ... of E and yet ... fail to cause ... E . Here probability increase is not sufficient for actual ... causation. How do we know that the envisioned case is not also one of these cases, in which C_1 fails to cause E ... and E is instead caused by C_2 [alone]? (Woodward 1994, p. 366; with minor changes in notation. Woodward notes that the example is originally due to Humphreys 1989, p. 15)

The potentially troublesome implication—that both carcinogens caused the stomach tumor—is not unique to Humphreys’s theory, but it is shared by all of the probabilistic approaches to singular causation mentioned above. Note that the assumption of indeterminism is essential to the example. If, given the relevant back-

ground conditions, the presence of diazonium metabolite increases the probability of stomach tumor to 1, then the presence of gyromitrin cannot further raise the probability.

Is Woodward's scenario coherent? Could it be that only one of the carcinogens actually causes the stomach tumor? Consider two potential analogues.

Scenario 2 Maria and Lucinda are neighbors. Every week, each of them buys a one-dollar ticket for the state lottery, which pays one million dollars as the grand prize. Due to the diminishing marginal utility of money, Lucinda and Maria decide that they would each rather have two chances at half a million dollars than one chance at a million, and so they decide to pool their resources. Each of them buys a ticket as before, but under the agreement that the prize money will be shared if *either* of them wins. In fact, Maria's ticket wins, and the neighbors split the million-dollar prize.

Scenario 3 Desmond and James decide to do something similar. They walk to the store together, slap their dollars down simultaneously, and purchase two tickets, each registered in both names. In fact, one of those tickets wins, and the two men share a million dollars.

In each scenario, how might we address the question: whose dollar purchased the winning ticket? Although the two women shared the money as agreed, and although each ticket increased the chance that the prize would be won, there is still a fact of the matter as to *whose* dollar had been converted to one million: It was Maria's. In scenario 3, however, there is no such fact. Although only one of the *tickets* actually won, there is no fact of the matter about whose dollar had actually purchased this ticket.

Probabilistic theories of causation presuppose that indeterministic causation works on analogy with the third scenario. Various causes increase the probability of an effect by contributing to a "probability pool." Once the probability of the effect is determined, the dice are cast, and the event either occurs or it does not. The individual causes make no additional contribution to the outcome; they bring it about (or not) only via their contribution to the probability pool. David Lewis and Paul Humphreys both defend the analogy between scenario 1 and scenario 3. Humphreys (1989) discusses scenario 1 explicitly:

In the example at hand, the situation with both carcinogens present *is* different from the situation with only one—the chance is higher than with either alone because both chemicals have contributed to the value of that chance. And that is all there is. To think otherwise is to conceive of the example in terms of a deterministic image where the tumor was "entirely caused"

by the first chemical and the second chemical was thereby irrelevant. But that is not the situation. The second chemical is not irrelevant on this (or any other) occasion for it contributes to the chance on this occasion, as does the first chemical, and after they have done this, *nothing else causal happens*. It is . . . a matter of sheer chance whether the tumor occurs or not. (pp. 36–37)

Lewis rejects an analogous counterexample: “[T]he objection presupposes that the case must be of one kind or the other: either *E* definitely *would* have occurred without *C*₂, or it definitely would not have occurred without *C*₂. But I reject the presupposition that there are two different ways the world could be, giving us one definite counterfactual or the other . . .” (Lewis 1986, p. 180, with minor modifications in notation). Both writers articulate a picture of probabilistic causation in which causes bring about their effects only via their contributions to a probability pool.

At issue is not merely the correctness of the probabilistic approach to causation—although that is important in itself—but how we are to conceive of indeterministic causation in general. This has implications that extend beyond metaphysics into risk analysis and tort law. Suppose, for example, that a pregnant woman is exposed to a particular teratogen, which increases the probability that her child will possess a particular kind of birth defect (which may also occur in the absence of the teratogen). Assume, moreover, that these are objective probabilities and do not merely reflect our ignorance about background conditions. If the birth defect occurs, is there a residual issue of whether it was caused by the teratogen, and hence of whether the woman may be entitled to receive payment for damages? American tort law answers in the affirmative, but it is plausible that this answer is founded on a denial of the assumption that the underlying probabilities are objective and not merely epistemic. If the relationship between exposure to the teratogen and the formation of birth defects is indeterministic, however, current jurisprudence may be deeply misguided on this issue. This critique of existing tort law is carried out in some detail by Parascandola (1996), although he assumes without argument that scenario 1 assimilates to scenario 3.

Although the intuitions of the opposing camps are clear, the case seems inconclusive. In particular, our immediate reaction to scenario 1 is to request more information. By what mechanism do the tumors form? How was the mouse exposed to the carcinogens? Which molecules of which substance made contact with which tissue cells? Which cells are the tumorous ones? And so on. In the absence of such further information, we have no clear intuitions about whether this case assimilates to scenario 2 or scenario 3. It will thus be helpful to examine some further examples that elicit clearer intuitions.

5 Three More Scenarios

Consider three more examples in which two potential causes are present and an effect occurs.

Scenario 4 Suppose that a source emits polarized photons where the angle of polarization depends on the physical configuration of the source. The emitted photons then strike a polarizer that is aligned so that horizontally polarized photons are transmitted. Initially, the source is configured so as to emit vertically polarized photons; any such photon will be absorbed by the polarizer. Now suppose that Juan and Jennifer simultaneously push the source, each pushing it with a force sufficient to rotate the apparatus through 30° . The source emits a photon, which now has a probability of $\sin^2(60^\circ) = 0.75$ of being transmitted by the polarizer, and the photon is in fact transmitted.

In this scenario, each push increased the probability that the photon would be transmitted. It seems clearly wrongheaded to ask whether the transmission was *really* due to Juan's push rather than Jennifer's push. Rather, this case seems to fit the model of probabilistic causality described by Humphreys and Lewis: Each push increased the probability of the effect (the transmission of the photon), and then it was simply a matter of chance. There are no causal facts of the matter extending beyond the probabilistic contribution made by each. This is a very simple case in which our intuitions are not clouded by tacit deterministic assumptions. If this case is typical of genuine indeterministic causation, then the sort of case envisaged by Woodward would indeed be incoherent.

Scenario 5 also supports this conclusion:

Scenario 5 A series of identical ropes are tied to a wall, and it is determined that when either Gene or Pat pulls on a rope as hard as possible for ten seconds, there is a fifty percent chance that the rope will break. When both pull on a rope for ten seconds, there is a seventy-five percent chance that it will break. On a given occasion, both pull and the rope does break. (Woodward, personal communication.)

Once again, our intuition is that it would be wrongheaded to say that on this occasion the breaking of the rope is due to the pull of Gene rather than Pat or vice versa. This example is less esoteric than the previous one, although we currently have no principled reason to think that in a scenario such as this, the probabilities would be irreducible and have this structure. This is a common problem in discussions of indeterministic causation: The only examples in which we have principled reasons for thinking that there are irreducible probabilities having a certain

structure involve microphysics, where our normal causal intuitions begin to lose their moorings.

Scenario 6 spells trouble.

Scenario 6 Suppose that two gunmen are shooting at a Ming vase. Each one has a fifty percent chance of hitting the vase, and each one shoots independently, so the probability that the vase shatters is 0.75. (For simplicity, we will ignore the possibility that the vase might survive a bullet hit.) As it happens, the first gunman's shot hits the vase, but the second gunman misses.

In this example, the second gunman shot at the vase, his shooting increased the probability that the vase would shatter (from 0.5 to 0.75), and the vase did in fact shatter. Nonetheless, it seems clear that we should not say that the second gunman's shot *caused* the vase to shatter. So here we have an apparent counterexample to probabilistic theories of causation: The second shot increased the probability that the vase would shatter, but it did not cause the vase to shatter. This example has the same structure as Woodward's, but here it is *clear* that the effect is due to one of the probability-raising factors, and not to the other. The two gunshots do not simply contribute to the probability of the vase's shattering, after which "nothing else causal happens." Something else causal does happen: The bullet fired from the first gun strikes the vase, while the bullet from the second gun misses. Note that this example threatens only probabilistic accounts of *actual* causation. This example would not lead us to withdraw the causal tendency claim that shooting at vases causes them to shatter. Likewise, even if Woodward's intuitions about scenario 1 are correct, that would not lead us to deny that gyromitrin or diazonium metabolite tend to cause stomach tumors.

In the next three sections, I will discuss prospects for replying to this counterexample. The results will be seen to be inconclusive: although no line of response suffers a knockout blow, none is wholly persuasive, either. In particular, many of the responses to be considered must postulate further empirical facts, beyond those presented above. While we do not know for sure that those facts do not obtain, we readily judge that the second shooter did not cause the vase to shatter without asking whether such facts obtain. Moreover, the facts adduced are sufficiently esoteric that it is implausible that our judgment about this case tacitly presupposes them.

6 Causal Processes

The shattering of the vase can be traced back to the firing of the first shooter's gun via the bullet that was fired from his gun, whereas there is no analogous process

connecting the second shooter to the target. It is natural to point to this asymmetry in judging that the firing of one of the guns, but not the other, was a cause of the vase's shattering.

Spatiotemporal causal processes, such as the speeding bullets of scenario 6, play a central role in Salmon's (1984) theory of causation (see also Dowe 1992; Salmon 1994). It is not clear, however, how the appeal to causal processes can resolve our problem. There *are* processes connecting the second shooter to the vase: photons and sound waves, for example. Intuitively, these are not processes of the right sort. Unfortunately, causal process theories such as Salmon's do not have the resources to formulate the distinction between causal processes "of the right sort" and those "of the wrong sort"; at the very least, such a distinction will presuppose something like the probabilistic theory of causation that is at issue. Thus, for example, if we try to get around this problem by saying that there must be a connecting process that transmits a sufficient quantity of energy to shatter the vase, we must then address the question of what constitutes a sufficient amount of energy. The answer to this question will almost certainly have to be formulated in terms of counterfactuals or *ceteris paribus* conditional probabilities.⁵

In any event, there are reasons that mitigate against requiring that causation always involve connection via continuous processes. Suppose that a gunman fires at a human target, but a secret service agent jumps in front of the gunman's bullet, allowing herself to be hit by it. The agent's brave action caused the target to remain alive, even though there are no processes connecting the agent and the target (except for irrelevant ones such as sound waves and photons).⁶ A probabilistic theory of causation thus cannot be saved by adding the simple requirement that causes be connected to their effects by continuous processes.

7 Precise Specification of Events

Consider the following scenario:

Scenario 7 Barney smokes, and he also spends a lot of time in the sun. These two proclivities are not connected; for example, Barney is not forced to go outside in order to smoke. Barney's smoking increases the probability that he will get lung cancer. By increasing his probability of getting lung cancer, Barney's smoking increases the overall probability that he will suffer from some form of cancer, and analogously for his exposure to the sun. In fact, Barney develops skin cancer. *A fortiori*, Barney develops cancer of some form or other.

Since Barney's smoking increased the probability that he would develop cancer of some sort, and Barney did indeed develop cancer, probabilistic theories of causation would seem to rule that his smoking caused him to develop cancer. But this seems clearly wrong: Barney's cancer is *skin* cancer, not lung cancer, so it was not caused by his smoking. Although both smoking and sun exposure increased the probability that he would develop cancer, only Barney's exposure to the sun actually caused him to get cancer.

Few philosophers would be troubled by this *prima facie* counterexample to probabilistic theories of causation. Our first reaction is that such theories yield the wrong verdict only because we have described the effect—Barney's getting cancer—imprecisely. There is a more detailed level of description—Barney's contracting skin cancer—where probabilistic theories render the correct verdict: Barney's smoking did not raise the probability of skin cancer, and hence it did not cause Barney's skin cancer.

Let us formulate the general principle that seems to be at work here:

(PSE) Suppose that on a particular occasion, events occur that instantiate types *C* and *E*. Even if, relative to the relevant background conditions, *C* raises the probability of *E*, if there are more precise specifications of the events in question, *C'* and *E'*, such that *C'* does not raise the probability of *E'*, then we should not say that *C* causes *E*.

How, precisely, one ought to formulate this principle of the precise specification of events will depend on one's theory of events—I will leave it to the reader to reformulate (PSE) in accord with her favorite theory of events.

Perhaps appeal to (PSE) will enable us to avoid the counterexample of the second shooter. Although the firing of the second shot increased the probability that the vase would shatter, there may exist a more precise specification of the outcome (say, involving the shattering into very small pieces of a certain region of the vase directly facing the first shooter) such that the second shot did not increase *its* probability.⁷ But although (PSE) seems to provide a natural diagnosis of scenario 7, it does not seem to get at our reasons for saying that the first gunman's shot caused the vase to shatter, while the second gunman's shot did not. Moving from the description "Barney's cancer" to the more specific "Barney's skin cancer" seems wholly natural, whereas the move from "the vase's shattering" to the more detailed level of description seems *ad hoc*.

Moreover, the application of (PSE) in this case involves an empirical presupposition: that there is some level of description such that the second gunman's shot made

no difference whatsoever to the probability that the vase would shatter in just that way. Perhaps the second gunman's shot did increase the probability, ever so slightly, that the vase would shatter *in precisely the manner that it did*. If so, (PSE) will be of no avail. If (PSE) succeeds at all, then, it will succeed because of empirical facts we have no (current) warrant for accepting, and which appear to play no role in our straightforward causal judgments about scenario 6.

8 Ebb and Flow in the Probability Pool

There is one regard in which scenarios 4 and 6 are significantly different. In scenario 4, both Juan and Jennifer push the polarizer, yielding a 0.75 probability that the photon will be transmitted. Nothing further happens to affect the probability of transmission once the position of the polarizer is set. This fits nicely with the probability pool conception of indeterministic causation: The two pushes determine the probability of transmission, and nothing else causal happens—transmission follows by sheer chance. Time is not literally discrete in this scenario, but it might as well be. It is implausible that scenario 6 works in this way: We do not imagine that the two guns fire, fixing the probability of shattering at 0.75 until the very moment of shattering. Rather, we imagine that the actual trajectory followed by the bullets depends on random gusts of wind, fluctuations of air pressure, or perhaps even on indeterministic fluctuations intrinsic to the bullets themselves. The initial 0.5 probability that each shot will hit the target is an average over all of these possibilities. As the bullets speed through the air, the probability that each one will hit the vase does not remain frozen at 0.5, but changes in response to indeterministic events that occur after the guns were fired. The probability pool does not remain calm, but rises and falls as probability flows in and out.

In cases such as that described in scenario 6, there may be a vast number of intermediate events whose chance occurrence will affect the probability of the outcome. Perhaps a probabilistic theory of causation that is sensitive to the ebb and flow of the probability of shattering after the two shots are fired could yield the right verdict about scenario 6. Such theories have been offered by Kwart (1997, *inter alia*), Menzies (1989a, but repudiated in 1996), and Eells (1991, chapter 6). I will not describe each of these theories in detail, but will rather focus on the features of the probability pool that they might exploit in dealing with the case of the two gunmen.

Suppose that there is some time t before the shattering of the vase, such that the trajectory of the bullet from second gun at t gives it no chance of hitting the vase: Although the bullet from the second gun emerged from the barrel with a probability

of 0.5 of hitting the vase, by time t , random changes in the bullet's trajectory have put it on a path that will definitely steer clear of the vase. Then, the probability at t that the vase will shatter is exactly the same as it would have been had the second gunman never fired. Under this supposition, the accounts of Menzies (1989a) and Eells (1991) correctly rule that second gunman's shot did not cause the vase to shatter. I will not defend these claims in detail here, but will instead focus on a simpler proposal suggested by Hall (personal communication). Hall recommends a minor modification of Lewis's counterfactual theory of causation: Instead of requiring that causes raise the probability of their effects *at the time the cause occurs*, we must require that causes raise the probabilities of their effects *shortly before the time at which the effect occurs*. In the case at hand, the probability at time t that the vase will shatter would have been no different if the second gunman had never fired; hence Hall's proposal rules that the second gunman's shot is not a cause of the shattering.

Another possibility emerges if instead we focus on the trajectory of the *first* bullet. Suppose now that there is some time t at which the bullet fired from the first gun is determined to hit the vase. If we hold fixed the trajectory of this bullet at time t , then the firing of the second gun no longer makes a difference to the probability of shattering: The probability is 1 either way. This move is not ad hoc: The trajectory of the first bullet at time t is a cause of shattering, and so should (according to many probabilistic theories of causation) be held fixed when evaluating the probabilistic relevance of the second gunman's shot. If the trajectory of the first bullet at time t is held fixed, the second gunman's shot no longer raises the probability of shattering; hence, a sufficiently refined probabilistic theory of causation correctly rules that the second shot is not a cause of shattering.

While these approaches hold out some promise for handling the case of the second shooter, each of them requires assumptions about the probabilistic details that were not included in the original description of the scenario. If the probability that the first bullet will miss the vase or that the second bullet will hit remains positive right up until the time of shattering, the proposed resolutions will not succeed.⁸ Although the probabilistic assumptions discussed above are not altogether implausible, our ignorance about such matters of probabilistic detail surely outstrips our ignorance of the causal relations in scenario 6. Our intuition that the first shot caused the shattering and the second did not does not depend on such probabilistic details. It is particularly troubling that the proposals discussed in sections 7 and 8 presuppose that certain probabilities take on values of 0 and 1. One of the motivations for developing probabilistic theories of causation was the thought that the world, for all we know,

might be indeterministic. It might, for all we know, be *thoroughly* indeterministic; perhaps everything is possible, no matter how improbable. Surely we should want our probabilistic theories of causation to apply in such a world as well.

9 The Last Scenario

I present for consideration a final scenario.

Scenario 8 A box contains two carbon-14 atoms, and a very sensitive detector is placed near the box. The presence of each atom increases the probability that a decay event will be detected during some given time interval. In fact, a decay event is detected, and one of the atoms in the box is now a nitrogen atom.

Intuitively it seems that it was the presence of one atom (the one that is now a nitrogen atom), and not the other, that caused the detection event.⁹ Yet in this scenario, *none* of the maneuvers canvased in sections 6 through 8 seems possible. Is this, at long last, a clean counterexample to probabilistic theories of causation?

We must proceed with caution here, for we are within the realm of the notorious measurement problem of quantum mechanics. There are many interpretations of what is going on in this example, no one of them universally accepted. As it turns out, those interpretations that lend themselves to the natural conclusion that it is the presence of one atom and not the other that caused the detection event tend not to be the ones that close the door on the various maneuvers described above. Suppose we understand the decay event in terms of a realist, or hidden-variable interpretation (such as that of Bohm 1952). In this case, one of the atoms decays at some time prior to the detection event, and a decay particle travels to the detector. There *is* a spatio-temporally contiguous process (the decay particle) connecting the atom to the detector, and there *is* a more detailed specification of the detection event available (one including hidden variables).

Suppose, by contrast, that we understand this scenario in terms of a more traditional Copenhagen-style interpretation. The atoms enter into superpositions of decayed and undecayed states, thus bringing about a certain probability that a measurement will reveal that a decay event has occurred. When the measurement is performed, the superposition collapses and a decay is in fact detected. On this picture, the detection of the decay particle is not a distinct event from the decay event itself. As an indeterministic result of that decay/detection event, one of the atoms is now a nitrogen atom. The decay of that atom is better thought of as a part of the measurement interaction than as a cause of it.

This is not an exhaustive partition of possible interpretations of scenario 8, of course; but it should suffice to convince us that this is the wrong place to look for a trouble-free counterexample to probabilistic theories of causation.

10 Conclusion

Note that our question is not one of epistemology—how do we know in cases like scenario 6 which potential cause is an actual cause? We do know—otherwise the counterexamples would have no force. In scenario 6, we know that the second gunshot did not cause the vase to shatter because the two events were not connected by an appropriate type of spatiotemporal process. In scenario 7, we know that Barney's smoking did not cause his cancer because Barney developed skin cancer, and that is not the sort of cancer that smoking causes. In scenario 8 (assuming some realistic interpretation of atomic decay), we know that the presence of one atom but not the other caused the detection event, because atoms cause detection events by decaying, and atoms that decay are transformed from an atom of one type to another. In each case, there is a *marker* that distinguishes the genuine cause from the spurious probability-raiser.¹⁰ But as metaphysicians, we are interested in providing a general theory of causation. The markers described above are heterogeneous in nature: in scenario 6, the marker resides in events intermediate between the putative cause and its putative effect; in scenario 7, it is the nature of the effect itself that marks the actual cause; in scenario 8, a separate effect (the transformation of the decaying atom) marks the actual cause of the detection event. For a theory of causation to exploit these markers, something more must be said about what they have in common in virtue of which they are causal markers. It will not do to simply say that the actual cause is the one that is marked as such.

It might be tempting to conclude that probabilistic theories of actual causation must ultimately fail. But there are still reasons for resisting this conclusion. The success of probabilistic theories of causal tendencies, together with the presumed connection between causal tendencies and actual causation, should lead us to resist the conclusion that actual causation has nothing to do with probability-raising. Even if we accept the example of the two gunmen as a direct counterexample to probabilistic theories of actual causation, scenarios 4 and 5 suggest that there are cases in which the probability pool model articulated by Lewis and Humphreys is appropriate. It will still be necessary, then, to map out the boundaries of this form of indeterministic causation, where probabilistic theories of causation directly apply. The counterexample of the second shooter, therefore, is not offered as a call to abandon prob-

abilistic theories of actual causation, but rather as a call for more work to be done. Whatever the result, our picture of causation in an indeterministic world can only come into sharper focus.

Acknowledgments

For helpful comments and suggestions, thanks go to Joseph Berkovitz, Phil Dowe, Ellery Eells, Dan Hausman, Paul Humphreys, Igal Kwart, Laurie Paul, Jon Schaffer, and Jim Woodward.

Notes

1. The condition described is only sufficient, according to Lewis. Since we will be primarily concerned with counterexamples to the sufficiency of probability-raising for causation, I will not present the more complex, weakened requirement that Lewis takes to be both necessary and sufficient.
2. Note, however, that not all of the probabilistic theories of causation mentioned above are intended by their authors as theories of actual causation.
3. Menzies does also mention problems due to “randomly occurring events” (1996, p. 100, n. 24), but he does not state clearly what sort of problems he has in mind.
4. Woodward notes that a similar counterexample is possible without the presence of C_2 , if there is some probability that an event of type E will occur spontaneously.
5. For a more detailed critique of causal process theories of causation, see my (1995b).
6. A similar example is described in Otte (1986); see also my (1998) and Hall, “Two Concepts of Causation” (chapter 9 in this volume).
7. Note that this proposal differs slightly from a similar proposal rejected by Lewis (1986b, pp. 204–205). That proposal was to move to a more precise specification of events in order admit the existence of certain causal relations otherwise denied by the theory; (PSE) aims to reject the existence of causal relations otherwise admitted by the theory.
8. If the probabilities remain positive but approach zero in the limit as the actual time of shattering nears, then suitable reformulations of the proposals discussed above may succeed. I will not explore this possibility further here.
9. Schaffer (2000a) presents an example with a similar structure.
10. This line of thought has been pressed on me independently by Joseph Berkovitz and Jonathan Schaffer. See also Hausman (1998), section 13.2.